[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Chaos Communication Congress
Berlin - bcc, 27.-30.12.2008
http://events.ccc.de/congress/2008/

.....

Proceedings [redacted] Datenp[redacted] Solar-powering [redacted] open source PLC tool[redacted] 202c StGB S[redacted] Smart Card [redacted] Sys[redacted]ng hackerspaces[redacted] Cyborgs and [redacted] Terrorist All-Stars [redacted] Staat [redacted] Virenprogrammierer? [redacted] and Georgia? D[redacted] digitale Intimsphäre [redacted] Reverse Engineering [redacted] Mühsams Tagebücher [redacted] Hacking the iPhone [redacted] Beyond Asimov [redacted] MSP430 BSL [redacted] Cold Boot [redacted] Locating Mobile Phones [redacted] European security [redacted] DNS [redacted] coreboot [redacted] Garage Doors [redacted] Gödel Jahres[redacted] [redacted]-Privacy [redacted] Climate [redacted] Disk-Encryption [redacted] Internet Applications [redacted] Hacking [redacted] Atmosphere [redacted] Symbian [redacted] [redacted] PHP [redacted] Music [redacted] Handschellen [redacted] smart-phone [redacted] MICA*-based wireless [redacted] TCP Denial [redacted] [redacted] Short Attention [redacted] Security [redacted] Trust Situation [redacted] [redacted] Swarm Robotics [redacted] Your Life [redacted] All your base[redacted] Banking Malware [redacted] Infinite Library [redacted] Wii [redacted] Fail Tricks: [redacted] Fnord News [redacted] Blinkenlights [redacted] a[redacted] Holodeck! [redacted] Soviet Unterz[redacted] your own GSM [redacted] Spuren [redacted] eVoting [redacted] TI EZ430U [redacted] Brother [redacted] Neusprech [redacted] Privacy [redacted] semantic web [redacted] stream cipher [redacted] Commodore 64 [redacted] Hacking [redacted] Botnets [redacted] anonymity [redacted] Tor Squeezing Attack [redacted] phones [redacted] Anonymous VPN [redacted] SWF [redacted] Attacks with Office Docu[redacted] Personalausweis Cisco IOS [redacted] Revolution [redacted] Wikileaks [redacted] Jeopardy [redacted] got owned [redacted] [redacted] technology sucks [redacted] Privacy [redacted] Quadrature du Net [redacted] RNG in [redacted] OpenSSL package Crafting a[redacted] theoretical possible [redacted] social contacts [redacted] [redacted] RFID [redacted] Pflanzenhacken [redacted] Nightmares [redacted]

Peter Fnord

Für Wolfgang.

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Chaos Communication Congress
Berlin - bcc, 27.-30.12.2008
http://events.ccc.de/congress/2008/

.....

.....
Proceedings .
.....

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Proceedings of the 25th Chaos Communication Congress
December, 27th - 30th 2008, Berlin Congress Center, Alexanderplatz.

25C3: An event of the Chaos Computer Club.
https://events.ccc.de/congress/2008/

Cover:	Evelyn Schubert
Data-Gardening:	Sven Klemm
Editor:	Matthias 'wetterfrosch' Mehldau
Publisher:	Art d'Ameublement Marktstraße 18 in 33602 Bielefeld
Vertrieb:	FoeBuD e.V. Unterstützungsbedarf Marktstraße 18 in 33602 Bielefeld, https://shop.foebud.org/
Font-Family:	Myriad Pro
ISBN:	978-3-934636-06-4
ISSN:	1867-8556

Program planning
under the patronage of the Wau Holland Foundation.

1st edition, 400 copies

Last update:	December, 15th 2008
Printer:	Druckerei Wollenhaupt Unter dem Felsenkeller 30 in 37247 Großalmerode
Paper:	Printed on FSC-certified paper by a FSC-certified printer.
Environment:	Climate-neutral print. The for the production this book emitted carbon-dioxide was compensated in cooperation with the natureOffice.de-program which supports facilities for renewable energy in developing countries. Strömchen!

W
HOLLAND
STIFTUNG

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Day 2008-12-27

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

18:30 CET

Das Grundrecht auf digitale Intimsphäre ... p. 31
Festplattenbeschlagnahme in neuem Licht
Saal 1: Society with Ulf Buermeyer, 46halbe

Chip Reverse Engineering ... p. 32
Saal 2: Hacking with Karsten Nohl, starbug ... see paper on p. 155

Erich Mühsams Tagebücher in der Festungshaft ... p. 33
Ein Idylle aus der Analogsteinzeit der Überwachung
Saal 3: Society with Johannes Ullmaier

20:30 CET

Hacking the iPhone ... p. 34
Pwning Apple's Mobile Internet Device
Saal 1: Hacking with pytey, MuscleNerd, planetbeing

Beyond Asimov - Laws for Robots ... p. 35
Developing rules for autonomous systems
Saal 2: Society with Frank Rieger

Cracking the MSP430 BSL ... p. 36
Part Two ... see paper on p. 165
Saal 3: Hacking with Travis Goodspeed

21:45 CET

Advanced memory forensics: The Cold Boot Attacks ... p. 37
Recovering keys and other secrets after power off ... see paper on p. 133
Saal 1: Hacking with Jake

Locating Mobile Phones using SS7 ... p. 38
Saal 2: Hacking with Tobias Engel

Collapsing the European security architecture ... p. 39
More security-critical behaviour in Europe! ... see paper on p. 263
Saal 3: Society with Gipfelsoli

23:00 CET

Why were we so vulnerable to the DNS vulnerability? ... p. 46
Saal 1: Hacking with Effugas

coreboot: Beyond The Final Frontier ... p. 47
Open source BIOS replacement with a radical approach to boot.
Saal 2: Hacking with Peter Stuge

Messing Around with Garage Doors ... p. 48
Breaking Remote Keyless Entry Systems with Power Analysis
Saal 3: Hacking with Timo Kasper, Thomas Eisenbarth

24:00 CET

Kurt Gödel – I do not fit into this century ... p. 49
Ein audiovisuelles Live-Feature
Saal 1: Culture with 46halbe, Marcus Richter, Ina Kwasniewski, Kai Kittler

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Day 2008-12-28

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Day 2008-12-29

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

16:00 CET

17:15 CET

18:30 CET

20:30 CET

21:45 CET

23:00 CET

Hacker Jeopardy ... p. 108
Die ultimative Hacker-Quizshow
Saal 1: Community with Stefan 'Sec' Zehl, Ray

We got owned by the (rhymes-with-unease) and didn't even get a lessons learned ... p. 109
life with targeted attacks
Saal 2: Hacking with jf

Day 2008-12-30

11:30 CET

Why technology sucks ... p. 112
If technology is the solution, politicians are the problem
Saal 1: Community with Brenno de Winter

Lightning Talks 3 ... p. 113
5 mins of fame
Saal 2: Community with SvenG

The Privacy Workshop Project ... p. 114
Enhancing the value of privacy in todays students view
Saal 3: Society with Christoph Brüning, Kai Schubert

12:45 CET

La Quadrature du Net - Campaigning on Telecoms Package ... p. 116
Pan-european activism for patching a 'pirated' law ... see paper on p. 269
Saal 2: Society with Jérémie Zimmermann, Markus Beckedahl

Predictable RNG in the vulnerable Debian OpenSSL package ... p. 118
the What and the How
Saal 3: Hacking with Luciano Bello

14:00 CET

Crafting and Hacking: Separated at Birth ... p. 119
Saal 3: Culture with Kellbot

15:15 CET

Making the theoretical possible ... p. 120
Attacking a critical piece of Internet infrastructure
Saal 1: Hacking with Jake, Alexander Sotirov

Mining social contacts with active RFID ... p. 121
Saal 2: Hacking with Ciro Cattuto

Pflanzenhacken ... p. 123
Züchten 2.0
Saal 3: Making with paul

16:30 CET

Security Nightmares 2009 ... p. 124
Oder: worüber wir nächstes Jahr lachen werden
Saal 1: Hacking with Ron, Frank Rieger

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Papers

Papers

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Chaos Communication Congress
Berlin - bcc, 27.-30.12.2008
http://events.ccc.de/congress/2008/

.....

.....
Lectures .
.....

Day 1

2008-12-27 | 11:30 CET | 01:00 h | Saal 1 | lecture | Society

Datenpannen

Forderungen nach dem Jahr der Datenverbrechen

Wer nichts zu verbergen hat, hat nichts zu befürchten? Die zuständigen Mitarbeiter halten sich strikt an das Gesetz? Überwachung hat für die Betroffenen keine negativen Folgen? Im Jahr 2008 sind diese Irrtümer so häufig widerlegt worden wie noch nie: Datenskandale bei LIDL, Telekom und dutzenden anderen, per Internet zugängliche Meldedaten, Massenverkauf von Bank- und Telefondaten – eine Liste ohne Ende im Datenskandaljahr 2008.

Wir nehmen die wichtigsten deutschen Datenskandale des Jahres unter die Lupe. Was war die Ursache? Welche Rechte habe ich als (möglicherweise) Betroffener? Und wie kann man das in Zukunft verhindern?

Damit ist es aber nicht getan, wir wollen gleichzeitig die Forderungen formulieren und diskutieren, die aus diesen Vorfällen folgen. Mit Schäubles neuem kleinen Datenschutzgesetz-Update wird sich jedenfalls nichts grundlegend ändern, daher ist es Zeit, unsere Vorstellungen für die Zukunft des Datenschutzes zu artikulieren.

> http://events.ccc.de/congress/2008/Fahrplan/events/2814.en.html

46halbe

Patrick Breyer

2008-12-27 | 11:30 CET | 01:00 h | Saal 2 | lecture | Making | See paper on p. 225!

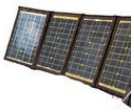

Solar-powering your Geek Gear

Alternative and mobile power for all your little toys

This talk will show you how to solar-power your laptop, PDA, cell phone, portable fridge or almost any other small device. Topics discussed include choosing the right solar panel, using (or not using) a voltage regulator, buffering the energy, some real applications as well as instructions on how to build a small and simple device to measure your power and energy savings.

Do you want to use your laptop in the garden or in the park without needing to pull long cords? Need to recharge your cell phone, PDA or camera in the wilderness? Are you just curious about solar energy or just want to keep your drinks cool on a hot summer day? Well, then you should attend this lecture!

Contents of the lecture:
• Motivation
• Decide what you want to have powered
• Choosing the appropriate solar panel
• Connectors, adapters, plugs
• The universal Buck/Boost voltage regulator
• Building your own device to measure voltage, current, power and energy
• Applications

> http://events.ccc.de/congress/2008/Fahrplan/events/2904.en.html

script

2008-12-27 | 12:45 CET | 01:00 h | Saal 2 | lecture | Community

U23

The Hackerspace's Junior Academy

Organize and operate a workshop for young people. Show them how your hackerspace works. Gain their attraction in having fun with hardware, electronics, microprocessors, software or hacking. Become known to new persons. Create networks of brains for new, cool projects. Let them experience the amazing power of teamwork!

In 2002, some people at the Chaos Computer Club Cologne discussed, how they could attract young people, especially students and pupils, to the ideas and lifestyle of a hacker, and gain new members. The result of this discussion was the concept for a project directed at young nerds and geeks, featuring a challenge which is only solvable as a group. This idea turned out to be so successful, that up to now there have been six recurrences.

The talk will explain the main design patterns which evolved in this six year period. We will introduce our motivation and goals for this project, and present the patterns for preparation, implementation and review.

We will save some time at the end for a short Q&A session, and fd0 is available in the hardware hacking room (in the basement) for a chat.

http://koeln.ccc.de/u23 *U23 at CCC Cologne*

> http://events.ccc.de/congress/2008/Fahrplan/events/2827.en.html

fd0

Lars Weiler

red_hood

2008-12-27 | 12:45 CET | 01:00 h | Saal 3 | lecture | Hacking

FAIFA: A first open source PLC tool

PowerLineCommunications has now their open source tool

PLC (PowerLineCommunications) had been widely used currently for the in-home LANs and for Internet access over PowerLineCommunications based on the market standard called HomePlug. Electricity is a great medium to transport data over existing cables in-home and outdoor but gives the network an old-school flavor of the behaviour of the hub where all stations share the medium. In this lecture, we present the freshly released FAIFA open source software that can be used to audit the security of PLC networks and script some flawnesses of the PLC devices.

PLC will definitely be one of the main LANs technology for in-buildings, in-home and collectivities IP connectivities in developed and undeveloped countries. PLC describes the technology used to developp MAC layer networks over existing power cables (110/220V - 50/60Hz) and TV cables in-building, in-homes and over public electrical networks.

This talk will describe the FAIFA tool and the technical overview of the current PowerLineCommunications technologies by outlining the following content:

- Overview of the PLC Networks
  - Introduction to the PLC, brief history of the technology
  - Technologies underlying the PLC (PHY and MAC layer)
  - Current status of the technology and security issues
- Description of the HomePlug AV standard
  - Different generations of chips and vendors
  - Possible Hacking targets (chip's architecture, on-chip system, managment of the chip)
  - PLC flooding, managment ARP spoofing, RAM dumping
- Description and demo of the FAIFA tool
  - Full description of the FAIFA features
  - Hacking audit using FAIFA
  - Call for contributions and developers

http://open-plc.org

> http://events.ccc.de/congress/2008/Fahrplan/events/2901.en.html

Nicolas Thill

Florian florian@openwrt.org
OpenWrt developper and co-founder of OpenPattern SARL developing fully open hardware and software FPGA-based router.

Xavier Carcelle

2008-12-27 | 14:00 CET | 01:00 h | Saal 1 | lecture | Hacking

Der Hackerparagraph 202c StGB

Bestandsaufnahme und Auswirkungen

Es wird Zeit, dass wir mal über die Dinge sprechen, die wir seit dem Inkrafttreten des Hackerparagraphen nicht mehr machen können.
Und die Dinge, bei denen wir uns nicht sicher sind, ob wir sie machen können, und daher lieber sein lassen.

Wir haben als Auswirkungen des Hackertoolverbots diverse üble Dinge vorhergesagt. Jetzt ist Zeit, mal zu gucken, was daraus geworden ist. Sind die Vorhersagen eingetroffen? Beschneiden wir Hacker uns in unserer Arbeit? Finden coole Projekte noch statt oder sind sie abgewandert?

> http://events.ccc.de/congress/2008/Fahrplan/events/3028.en.html

Felix von Leitner

lexi
Lexi Pimenidis is known for teaching offensive IT security classes.

Jan Münther

Jürgen Schmidt

2008-12-27 | 14:00 CET | 01:00 h | Saal 3 | lecture | Hacking

Security Failures in Smart Card Payment Systems

Tampering the Tamper-Proof

PIN entry devices (PED) are used in the Chip & PIN (EMV) system to process customers' card details and PINs in stores world-wide. Because of the highly sensitive information they handle, PEDs are subject to an extensive security evaluation procedure. We have demonstrated that the tamper protection of two popular PEDs can be easily circumvented with a paperclip, some basic technical skills, and off-the-shelf electronics.

PIN entry devices (PEDs) are critical security components in Chip & PIN (EMV) smartcard payment systems as they receive a customer's card and PIN. Their approval is subject to an extensive suite of evaluation and certification procedures. We have demonstrated that the tamper proofing of PEDs is unsatisfactory, as is the certification process.

This talk will discuss practical low-cost attacks on two certified, widely-deployed PEDs – the Ingenico i3300 and the Dione Xtreme. By tapping inadequately protected smartcard communications, an attacker with basic technical skills can expose card details and PINs, leaving cardholders open to fraud. The talk will describe the anti-tampering mechanisms of the two PEDs and show that, while the specific protection measures mostly work as intended, critical vulnerabilities arise because of the poor integration of cryptographic, physical and procedural protection.

These failures are important not only because they allow fraud to be committed, but also because of their affect on customer liability. As Chip & PIN was claimed to be foolproof, victims of fraud often find themselves accused of being negligent, or even complicit in the crime. The results of this work will help customers in this position argue that their losses should be refunded.

http://www.cl.cam.ac.uk/research/security/banking/ped/ Further information

> http://events.ccc.de/congress/2008/Fahrplan/events/2953.en.html

Steven J. Murdoch Steven.Murdoch@cl.cam.ac.uk
Steven J. Murdoch is a researcher in the Security Group of the Cambridge University Computer Laboratory, working on the Tor Project.

2008-12-27 | 16:00 CET | 01:00 h | Saal 1 | podium | Community

Building an international movement: hackerspaces.org

What we did so far. What will happen in the future.

We live in interesting times to build hacker spaces: physical spaces where hackers make things, inspired by European models, pop up everywhere. Whether you need inspiration to build your own hacker space or want an update on what happened in places like New York City, Washington D.C., San Francisco, or Vienna since last year: This international panel will provide you with insight.

We have come a long way since 2007. Looking at how things are done in Europe has inspired several new spaces in interesting corners of the world. What started as research at last year's Chaos Communication Camp and a talk on design principles for hacker spaces at 24C3 didn't stop there. Turning the theory of places for people, tools, and Club-Mate into an international movement, we have recently launched hackerspaces.org as a central hub for information exchange world wide. So far, the results are amazing: Spaces from almost all continents have joined us, and stories of success and inspiration can be told.

It doesn't take you very much to join: Four people can start a sustainable hacker space. There are few excuses left for not joining the global hacker space movement with a place of your own. This panel will cover building a hacker space, fab labs, co-working spaces, and other tech-oriented 'third spaces'.

> http://events.ccc.de/congress/2008/Fahrplan/events/2806.en.html

Nick Farr

Enki paul@boehm.org

Jens Ohlig

CCC activist for almost 2 decades. He likes words and code, and thus linguistics, literature, and computer sciences. According to Wikipedia criteria he is irrelevant, but thinks he is in excellent company.

Bre

Jake

2008-12-27 | 16:00 CET | 01:00 h | Saal 2 | lecture | Science | See paper on p. 235!

About Cyborgs and Gargoyles ...

State of the Art in Wearable Computing

In this talk I present the current state of wearable computing, computing as common and useful as clothes, focusing on activity recognition (the inference of the users current actions) using on-body sensors (accelerometers, gyroscopes and other modalities), explaining possibilities, dealing with challenges and limitations and presenting some perils.

I will also give some practical tips about the design, implementation and use of wearable computing systems from Thad Starner's Twiddler to the Wearable Labs QBIC from ETH.

Introduction

The visions of ubiquitous / pervasive computing more and more becomereality. Everyday, we are surrounded by a multitude of computing devices. However, as of today, they fulfill very limited functionality and often are more distractive than useful. In my talk I will present research efforts to eliminate this gap and to enable everyday technology to support us during real life without hinderance or annoyances focusing on wearable technology. Using portable devices that accompany us like clothes, we are supported by computing not relying on any kind of infrastructure and augmented spaces. I focus in my talk specifically on the state of the art in activity and context recognition.

Application scenarios

First, I will give some application scenarios from European Union Projects, we at the Embedded Systems Lab at the university of Passau are currently involved in. This includes WearIT@Work, a 500 Million Euro integrated project with over fifty partners. I will show short movies about some work we did at Skoda to support assembly line work and work involving maintenance work support for Zeiss technicians using Heads-up displays.
The next application scenario to tackle is Healthcare. Here we work closely with the Hospital in Steyr, Austria, supporting doctors doing their rounds using RFID technology and capacitve sensing. Another aspect is also health and lifestyle, an interesting personal hobby project involves trying to recongize Kung-Fu and Tai Chi gestures.

Enabling technologies

In this seciton of the talk, I will go into more details about what sensors and modalities can be used to detect which activities/ user states etc. From the pervasive accelerometers and gyroscopes integrated in a lot of gadgets (from mobile phones to entertainment consoles) over using a in-ear microphone to get chewing sounds and detect some kinds of food to a mobile phone that can detect the types of surfaces using simply vibration and audio fingerprinting. Other enabling technologies developed at ESL Passau include magnetic

2008-12-27 | 16:00 CET | 01:00 h | Saal 2 | lecture | Science | See paper on p. 235!

coils for distance measurements and fabric stretch sensors for muscle activity.
Most state of the art context and activity recognition techniques rely on a fixed number of sensors with known position and orientation. As part of an effort to facilitate wearable activity recognition using dynamically changing sets of sensors integrated in everyday appliances such as phones, PDAs, watches, headsets we try to tackle some key issues of such systems as detecting automatically on-body device position and developing displacement indifferent classification algorithms.

Practical tips

In this section, I describe some of the tools and gadgets that can currently be used to enhance us and support us in everyday life. I tackle the Twiddler (a one hand keyboard), the QBIC (a belt intergrated Linux PC from the ETH wearlab with whom we collaborate closely), yet also commercial systems like the oqo and the nokia 810.
I will also mention a context logger application developed by Tobias Franke, a student writing his diploma thesis at our institute. The application enables to save data from the iPhone 'sensors': audio, acceleration etc. to file and label it with some acitivity/ state label for later analysis.
The application is not finished yet. However, it will be free and opensourced. To the time of the congress, the first version will be available on the AppStore (hopefully ...).

http://esl.fim.uni-passau.de/ ESL Uni Passau website
http://opportunistic.de/ Effort to simplify activity recognition

> http://events.ccc.de/congress/2008/Fahrplan/events/2892.en.html

kai_ser
Kai is a n-th (n for an arbitrary large number) year phD. student at the Embedded Systems Lab at University of Passau, Supervisor Prof. Paul Lukowicz.

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-27 | 17:15 CET | 01:00 h | Saal 1 | lecture | Society

Terrorist All-Stars

Some cases of terrorism around the world that are not terrorist at all

After more than a year of mostly dealing with the terrorism investigation against my partner Andrej Holm, and the resulting total surveillance directed at him and our family, it has become more quiet lately for us. The investigation is *still* going on though.
In the course of my new preoccupation 'terrorism' I keep hearing about similarly absurd cases of such investigations. All different, but all with analogies. All hard to bear for those who are subjected tot them. The talk will introduce some cases and search for patterns in cases against 'terrorists' who are clearly not terrorists.

Some of the cases I'd like to present some details about:
Two researchers from Nottingham, UK, were arrested after one of them had downloaded (quite legally) an Al-Qaeda handbook and the other had helped him print it. The latter barely avoided being deported to Marocco as a result. -> freehicham.co.uk
Portugal has passed a new law against terrorism after 9/11. Next to several cases against right-wingers there is one investigation going on against anti-authoritarian political activists, who destroyed a genetically modified corn field, in broad daylight and with media accompanying them. There is video coverage of this. They are now awaiting trial. -> solimove.liveinfo.nl
Ten animal rights activists from Austria will face trial eventually after having spent more than three months in pre-trial detention. They are accused of forming a criminal organisation, based mainly on the assumption of 'conspiratorial behaviour' (they used encryption) and the fact that they active in the field of animal rights or animal protection -> antirep2008.lnxnt.org.
New Zealand also has a new 'Terrorism Suppression Act', and the first case was opened againt 17 people, partly Maori, partly white, who are accused of having violated the Arms Act. The terrorism charge has in the meantime been dropped, but pre trial hearings started were in September and opened in October -> october15thsolidarity.info
Along with these cases I will mention several German cases of so-called terrorism that were dropped by the prosecutor this year. They include one against alleged members of the 'militant group' (same charge as Andrej, but different people), one against the alleged members of the 'militant campaig against the G8 summit' (as was invented by the prosecutor) and a third against approx. ten young antifascist activists in a small town in northern Germany.

> http://events.ccc.de/congress/2008/Fahrplan/events/2991.en.html

Anne Roth annalist@riseup.net
Anne Roth lives in Berlin. She is a media activist, journalist, translator and mother of two. Since July '07 she's mostly known as the partner of Andrej Holm who was arrested at 7 in the morning in their apartment for being terrorist. Since October of that year she writes a web log about the inside view of an anti-terror investigation and increasingly about the War on Terror in general.

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-27 | 17:15 CET | 01:00 h | Saal 2 | lecture | Society

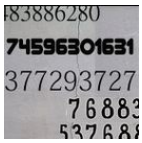

Der Staat als Virenprogrammierer?

Die Steueridentifikationsnummer als Gefahr der informationellen Selbstbestimmung

Seit August diesen Jahres verteilt das Bundeszentralamt für Steuern die neuen Steuernummern. Die elfstellige Ziffer ist der Traum aller Programmierer: Die Nummern sind eineindeutig, Verwechslungen ausgeschlossen - der ideale Index.

Auch wer noch nicht oder nicht mehr Steuern zahlen muss, bekommt eine Steuernummer; künftig werden sie mit der Geburt verteilt, bleiben bis zu zwanzig Jahre nach dem Tod gültig. Datenschützer fürchten, dass mit der Steuer-ID die letzten Hindernisse für den freien Datenaustausch zwischen staatlichen Behörden beseitigt werden sollen.

Welche Folgen wird die Verbreitung der Steuer-ID haben? Welchen Sinn, welche Realisierungschancen hat das Verbot der Bildung von Persönlichkeitsprofilen? Und warum ist die Zweckbindung von Daten so wichtig? Wo wird die Steuer-ID bereits gespeichert, wer will noch Zugriff darauf? Wie kann ich mich gegen Speicherung und Verwendung der Nummer wehren?

> http://events.ccc.de/congress/2008/Fahrplan/events/3039.en.html

Sven Lüders
Sven ist Geschäftsführer der Humanistischen Union.

2008-12-27 | 17:15 CET | 01:00 h | Saal 3 | lecture | Culture

Just Estonia and Georgia?

Global-scale Incident Response and Responders

Estonia and Georgia are just two examples of where global scale cooperation is required for handling security incidents on the Internet.

DDoS, fast spreading worm and 'CYBER WARFARE' are miniature examples of what the Internet faces every day. In this talk we will discuss how incidents are handled and specific case studies to illustrate it.

How is such large scale incident response handled? Who handles it?
Is that how it should be working?

What do YOU do if you need to initiate it, and should you just pray instead? :o)

> http://events.ccc.de/congress/2008/Fahrplan/events/2988.en.html

gadi

Gadi Evron is recognized for his work and leadership in Internet security operations and is arguably the world's top expert on botnets. He is also considered an expert on corporate security, counterespionage, and cybercrime (e-fraud and phishing). Previously, he was CISO at the Israeli government ISP (eGovernment project) and founded the Israeli Government CERT. He organizes and chairs worldwide conferences, vetted working groups and task forces. He authored two books on information security and is a frequent lecturer.

2008-12-27 | 18:30 CET | 01:00 h | Saal 1 | lecture | Society

Das Grundrecht auf digitale Intimsphäre

Festplattenbeschlagnahme in neuem Licht

Das Bundesverfassungsgericht hat uns anläßlich der Verfassungsbeschwerde gegen das nordrhein-westfälische Verfassungsschutzgesetz ein neues Grundrecht auf Gewährleistung der Vertraulichkeit und Integrität von informationstechnischen Systemen geschenkt. Damit wurden für den Einsatz des geplanten Bundestrojaners zwar genaue Regelungen getroffen, aber was ist eigentlich mit den tausenden Festplatten, die jedes Jahr in Deutschland beschlagnahmt werden?

Wir wollen den Wortlaut und den Sinngehalt des neuen Grundrechts und die Folgen daraus hinsichtlich der Auswertung von beschlagnahmten Festplatten analysieren. Müßte nicht der vielzitierte 'absolut geschützte Kernbereich der privaten Lebensgestaltung' auch beachtet werden, wenn digitale Speicher nach einer Beschlagnahme ausgewertet werden? Was steht dazu im Urteil zur Online-Durchsuchung und welche Änderungen sollten daraus in der Praxis folgen?

Um diese Fragen zu beantworten, werden wir zunächst den heutigen Alltag der Festplattenbeschlagnahme und -auswertung beschreiben. Wie eine grundrechtskonforme zukünftige Praxis aussehen kann, wollen wir dann skizzieren.

> http://events.ccc.de/congress/2008/Fahrplan/events/2923.en.html

Ulf Buermeyer

Ulf Buermeyer ist Richter der Großen Strafkammer 22 - Schwurgericht - des Landgerichts Berlin und Redakteur der Fachzeitschrift HRRS. Bis August 2008 war er als wissenschaftlicher Mitarbeiter an das Bundesverfassungsgericht abgeordnet.

46halbe

2008-12-27 | 18:30 CET | 01:00 h | Saal 2 | lecture | Hacking | See paper on p. 155!

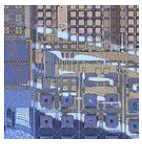

Chip Reverse Engineering

Cryptographic algorithms are often kept secret in the false belief that this provides security. To find and analyze these algorithms, we reverse-engineering the silicon chips that implement them.

With simple tools, we open the chips, take pictures, and analyze their internal structures. The talk provides all the details you need to start reversing chips yourself. Happy hacking!

> http://events.ccc.de/congress/2008/Fahrplan/events/2896.en.html

Karsten Nohl
Karsten is a hardware hacker and cryptographer. His academic research deals with privacy protection, while his hacking projects focus on cracking cryptographic hardware. This year, Karsten talked a lot about smart-card security and embedded cryptography at USENIX Security, BlackHat, CanSecWest, Toorcon, the HOPE conference, and other venues.

starbug

2008-12-27 | 18:30 CET | 01:00 h | Saal 3 | lecture | Society

Erich Mühsams Tagebücher in der Festungshaft

Ein Idylle aus der Analogsteinzeit der Überwachung

Während seiner Festungshaft (1920-1924) wurden dem Dichter und Anarchisten Erich Mühsam mehrfach die Tagebücher konfisziert, ausgewertet und (teils öffentlich) gegen ihn verwendet.

Der Vortrag schildert, welche absurden Bumerangwirkungen sich aus diesem Übergriff ergeben haben. Er fragt, inwieweit das, was im Rückblick als Entgleisung präfaschistischen Klassenjustizvollzugs erscheint, heute allgemeiner Standard zu werden droht und welche Abwehrstrategien sich aus der analogen Urgeschichte fruchtbar machen lassen.

> http://events.ccc.de/congress/2008/Fahrplan/events/2807.en.html

Johannes Ullmaier

Dr. Johannes Ullmaier ist Mitherausgeber der im Ventil Verlag erscheinenden Buchreihe »testcard – Beiträge zur Popgeschichte«, Dozent am Lehrstuhl für Neuere deutsche Literaturgeschichte an der Universität Mainz sowie als Lektor beim Suhrkamp Verlag tätig.

2008-12-27 | 20:30 CET | 01:00 h | Saal 1 | lecture | Hacking

Hacking the iPhone

Pwning Apple's Mobile Internet Device

Apple's iPhone has made a tremendous impact on the smartphone market and the public consciousness, but it has also highlighted their desire to carefully control the device with draconian restrictions. These restrictions prevent users from choosing to run third-party applications unauthorized by Apple and using the devices on carriers not approved by Apple.

Since its release, a tremendous amount of effort has been made to remove these restrictions for the benefit of the community. A year later, we have now learned much about its inner workings and have methods to circumvent these restrictions. This talk will summarize what we have learned about the internal architecture of the iPhone platform, its security, and the ways we have found to defeat these security measures.

http://www.iphone-dev.org *Our website/wiki*
http://blog.iphone-dev.org *Our blog*

> http://events.ccc.de/congress/2008/Fahrplan/events/2976.en.html

pytey

MuscleNerd

planetbeing

2008-12-27 | 20:30 CET | 01:00 h | Saal 2 | lecture | Society

Beyond Asimov - Laws for Robots

Developing rules for autonomous systems

Robotic systems become more and more autonomous, and telepresence develops very rapidly. But what happens if things go wrong? Who is responsible for that autonomous cleaning car murdering tourists? How can you identify the owner of that spy-drone filming you naked at the pool? This talk outlines some ideas to trigger a debate on how to deal with these problems, without stifling innovation and fun.

Asimovs three laws of robotics are the first thing that comes to mind when the 'how should robots be regulated' question comes up. However, with the current level of technology these 'laws' are irrelevant and can not be implemented. But we need other rules and laws to govern the use of autonomous and telepresence systems. Clear responsibilities need to be defined and enforced, without stifling innovation, development and fun.

The talk will suggest the development of a 'P2P TÜV' system for people and groups who build autonomous and telepresence systems before there are relevant official laws and regulations. The core idea is to think out and test practical ways to cope with the risks and uncertainties, so that there is a relevant body of experiences when the debates about official laws and rules begin. Building on the experience of the experimental aircraft movement, a P2P sanity and safety check system seems to be the right way to do this.

The second part of the talk will discuss possible regulation areas like laws, insurance rules, type approvals and number plates where experiences can be drawn from existing fields of technology regulation like steam engines and cars. Developing a position of the hobbyist and hacker community on robotic law may seem a bit early. But experience shows that technology development is fast and we need to come up with suggestions and ideas before a mad luddite mob does, or hordes of armed robocops roam the streets and skies.

> http://events.ccc.de/congress/2008/Fahrplan/events/2815.en.html

Frank Rieger

2008-12-27 | 20:30 CET | 01:00 h | Saal 3 | lecture | Hacking | See paper on p. 165!

Cracking the MSP430 BSL

Part Two

The Texas Instruments MSP430 low-power microcontroller is used in many medical, industrial, and consumer devices. When its JTAG fuse is blown, the device's firmware is kept private only a serial bootstrap loader (BSL), certain revisions of which are vulnerable to a side-channel timing analysis attack. This talk continues that from Black Hat USA by describing the speaker's adventures in creating a hardware device for exploiting this vulnerability.

While the previous part focused on the discovery of the timing vulnerability and its origin, this lecture will focus on the exploitation. Topics include a brief review of the vulnerability itself, PCB design and fabrication, the malicious stretching of timing in a bit-banged serial port, observation of timing differences on the order of a microsecond, and the hell of debugging such a device.

> http://events.ccc.de/congress/2008/Fahrplan/events/2839.en.html

Travis Goodspeed travis@utk.edu
Travis Goodspeed is a neighborly fellow from Knoxville in Southern Appalachia.

2008-12-27 | 21:45 CET | 01:00 h | Saal 1 | lecture | Hacking | See paper on p. 133!

Advanced memory forensics: The Cold Boot Attacks

Recovering keys and other secrets after power off

Contrary to popular assumption, DRAMs used in most modern computers retain their contents for seconds to minutes after power is lost, even at operating temperatures and even if removed from a motherboard. Although DRAMs become less reliable when they are not refreshed, they are not immediately erased, and their contents persist sufficiently for malicious (or forensic) acquisition of usable full-system memory images.

We show that this phenomenon limits the ability of an operating system to protect cryptographic key material from an attacker with physical access. We use cold reboots to mount attacks on popular disk encryption systems — BitLocker, FileVault, dm-crypt, and TrueCrypt — using no special devices or materials. We experimentally characterize the extent and predictability of memory remanence and report that remanence times can be increased dramatically with simple techniques. We offer new algorithms for finding cryptographic keys in memory images and for correcting errors caused by bit decay. Though we discuss several strategies for partially mitigating these risks, we know of no simple remedy that would eliminate them.

http://citp.princeton.edu/memory/
http://citp.princeton.edu/memory/code/
http://www.appelbaum.net

> http://events.ccc.de/congress/2008/Fahrplan/events/2922.en.html

Jake

2008-12-27 | 21:45 CET | 01:00 h | Saal 2 | lecture | Hacking

Locating Mobile Phones using SS7

You are used to your mobile phone number following you around the globe. But the same functionality that makes you reachable worldwide can also be used to track your whereabouts down to city-level – without you ever knowing about it.

This talk will explain what SS7 features are exploited for locating mobile phones, how the returned information has to be interpreted and what you can (and can't) do against being located that way without having to turn off your phone altogether.

> http://events.ccc.de/congress/2008/Fahrplan/events/2997.en.html

Tobias Engel

2008-12-27 | 21:45 CET | 01:00 h | Saal 3 | lecture | Society | See paper on p. 263!

Collapsing the European security architecture

More security-critical behaviour in Europe!

At the latest since 9/11, the EU took severe changes in their home affairs policy. New agreements and institutions were created to facilitate police networking (Europol, Frontex, CEPOL, new databases and their shared access). The european 'cross border crime fighting' has become an EU framework. Providing that this should help to win a 'war on terrorism', lots of the changes follow the US model of 'Homeland Security'. Risks' should be minimized by taking more and more 'proactive' measures and foresee possible 'threats'.

This includes the development of an 'Homeland Security Industrial Complex', whose budget is growing rapidly since 2001. The EU set up the research program 'FP7', that should help to find technical solutions for the discrete, but efficient surveillance and control. Both the research and the implementation of these new technologies are common projects of the security industry, military and police. This includes the usage of satellite pictures (whose resolution is now down to 50cm), drones (that are used in the EU already for catching migrants), geo-data/ mapping and several technics for border control. Satellites help to detect for example automatically deviant behaviour of vehicles (like migrants vessels).

Large changes are on their way to master the 'data tsunami' (as the EU ministers of interior call it). The problem for authorities is not anymore the gathering of data, but the processing. Streams are getting broader by transmitting real time data such as digital surveillance cameras (replacing the analog systems), biometrics, interception of communication data, command control units etc. The industry offers a variety of software to support the police to navigate within the immense flood of data. This software is sometimes pure investigative, used for datamining, but might be used also on risk profiling and foreseeing crimes (which is the 'evolution of crime fighting', as the developper proudly states). Police laws are to be changed for allowing the usage of software.

In the lecture we try to show the changes inside the EU towards a common politics of home affairs. Our examples will be their usage at large events like the FIFA 2006, G8 2007, Euro08. We follow the thesis, that the EU militarizes social conflicts in awaiting large changes in the 'global security architecture. Migration and climate change are seen as heavy 'security risks' for the EU.

Resistance against the described paradigm changes must leave the national levels. Like the campaigns against data retention, their implementation is decided by EU bodies - which are directed by the national ministers of interior. We would like to invite people around CCC to participate at a campaign focussing on EU ministers of interior, that will decide their next 5 years plan on european inner security in November 2009 in Stockholm, trying to intensificate the european panopticon (see link to first call below). To know what

2008-12-27 | 21:45 CET | 01:00 h | Saal 3 | lecture | Society | See paper on p. 261!

you can expect from the lecture, see one of our earlier presentations (see below).

Collapsing the European security architecture: For greater security-critical behaviour in Europe

Each protest enables us to draw conclusions of how to do things better next time. In the same way, we can draw conclusions from the mobilisation against the G8 summit 2007 in Heiligendamm on how to achieve successful and broad resistance. Apart from three large self-organised protest camps and an international infotour in the months leading up to the summit, there were attempts to have international exchanges and establish networks beyond Germany. The decision was made not to respond to the G8 climate debate but to frame the protests in terms of other self-determined topics the movement was focussing on: migration, antimilitarism and global agriculture.

Looking ahead to the 60th NATO anniversary in Strasbourg and Kehl and the G8 2009 in Italy, but also to the swedish EU Presidency 2009, this text takes up these points to propose a campaign against the new 'European Security Architecture'. We outline some developments in police cooperation on a European level and call for a kind of antirepression work that goes beyond a simple critique and a scandalising police violence, and that is coordinated on a European level. Such political antirepression work would have to take new forms of social control seriously as an integral reference point for radical movements.

No future for freedom

At the latest since September 11th 2001, not only the foreign policy coordinates of the European Union (EU) have changed. Under the motto 'Terror Comes Home', far-reaching changes in European Home Affairs, along with police operations towards a 'preventive security state' have been implemented. Whilst control of the external EU border has been stepped up with new technologies and cross-border cooperation, surveillance and control within the EU is also steadily increasing. Additionally, there are foreign military and police operations on behalf of the EU in so-called 'third countries'. The EU intends to be a model for a security complex that can be exported to other countries in the EU's capacity as a 'service provider'. These developments are not only directed at migrants and 'security-critical' behavior. They also offer a welcome opportunity to control the re-emerging alterglobalisation movement.

Since 1999, the EU has defined Europe as a 'space of freedom, security and law'. In future there will be more juridical and police cooperation in criminal and civil affairs. Home Affairs ministers dream of an EU ministry for Home Affairs. On the police level, EU bodies have received more competences, and new institutions and programmes have come into existence. In 2007, the so-called 'Future Group' met for the first time. This group is made up of the ministers of Home Affairs of the countries due to hold the EU presidency in the next four years. The EU commissioner for 'freedom, security and law' is also part of this group, along with the director of the 'Border Protection Agency' Frontex. The Future

2008-12-27 | 21:45 CET | 01:00 h | Saal 3 | lecture | Society | See paper on p. 261!

Group calls itself 'informal', but it has considerable influence on Home Affairs with respect to the EU Treaty as well the 2007 Lissabon negotiations. The foundation of the Future Group coincided with the EU presidency of Germany in 2007. Under the motto 'Living a secure Europe', the German Home Affairs Minister successfully pushed through a tightening of European internal policies[1].

Cross-border cooperation

Until now, cross-border police cooperation has only existed between some countries under the Pruem Treaty. This found its expression, for example, during the G8 summit 2003 when German police participated in an operation against demonstrators in Geneva with 500 police officers and five water canons. The Pruem Treaty was a test case and has subsequently been integrated in the 'legal framework of the EU'. Thus it now applies to all EU countries. All police departments will now have access to DNA and fingerprint databases as well as vehicle registration data. Access to information on 'terrorism suspects and travelling violent criminals' will be made easier in order to prevent travel or to 'quickly recognise and detain rioters'. For the European Football Championship in 2008, 2 000 German police officers have been ordered from Austria and Switzerland.

As an intersection for police cooperation, the competencies of Europol in The Hague are not restricted to gathering data and advising police forces of EU member states. An EU parliamentary decision in January 2008 meant that the 'European Police Office' became an EU agency for the 'coordination, organisation and implementation of investigative and operational measures'. The realms of responsibility have been extended to 'organised crime' and 'other forms of serious crime'. In future, access to the 'Europol Information System' will not require a 'liaison officer' anymore.

These 'liaison officers' are sent by the police forces of all member states to European control and decision-making bodies and are key figures in the policing of major events. Officially they have an 'advisory function'. In practice, they function as important nodes in informal police cooperation. They have access to all the databases of their home countries and can, for example during summit protests, provide information about different political groups. Liaison officers coordinate entry restrictions which led to 600 people being denied entry into Germany during the G8 2007, because they had been 'conspicuous during previous G8 summits'.

Europe – a space of surveillance and control

The cooperation between police and intelligence services is being expanded. In Germany, the Federal Criminal Investigation Office and the 'Verfassungsschutz' (Office for the Protection of the Constitution) recently moved to a 'Joint Terrorism Defence Centre', where they have separate offices but meet daily for joint briefings and share the canteen space. This cooperation led to the surveillance of the anti-g8-movement and the start of investigative operations under the premise of terrorism suspicions. German terrorism legislation allows for far-reaching interferences in people's privacy and allowed a record to

be taken of all mobile phone numbers present at a meeting of the radical left dissent!-network against the G8. As people affected by these operations have been able to access their files, it has come to light that these investigations were carried out by the police but initiated by the intelligence services.

Internet surveillance has increased across Europe. The German Ministry for Home Affairs has started a European-wide initiative to fight 'international terrorism', entitled 'check the web'. On March 8th 2007, Europol's 'information portal' went live. German police and secret services intend to cooperate with a joint 'internet monitoring and analysis project' in the future. Such 'internet surveillance centres' are planned across Europe. The intention is to partially automatise the monitoring of websites and subsequent archiving in police databases. New software scans the databases to find 'entities', which are conceptual analogies or connections between persons and objects ('semantic technologies'). The security industry is developing programmes that are able to search in different file formats. This way, text, audio, video and gps data can be analysed together. Prosecution agencies of various countries already use software that enables the 'prediction of crimes' as a result of data analysis. One company describes this process as an 'evolution in fighting crime'.

More police repression and law enforcement can also be observed in other countries of Europe. In Italy, several trials in relation to the G8 2001, as well as demonstrations against militarism and fascism concluded with sentences between 6 and 12 years for the accused. In other countries, police laws are being changed in order to give police more powers against 'security-critical behaviour'.

Radical changes have been made in Europe under Sarkozy and Berlusconi. In France, passengers who are the first to stand up to protest against a deportation on their flight risk being charged with ringleadership. New legislation in Italy has allocated 2,500 military troops for assistance in police operations to 'maintain public order'. The police intend to fingerprint any children of Roma origin found unaccompanied in the streets.

The new Austrian legislation on security police makes the racist control of migrants easier. The German Federal Police now have more competencies both for missions abroad and for domestic affairs, for example against political protests. EU member states implement European directives and 'harmonise' their national legislations, for example with respect to data retention. Telecommunication and internet providers now have to store data and hand it over to the police on request. This enables the police to reconstruct communications and create 'relational diagrams'. Protection from surveillance is increasingly restricted. The users of encryption software in Austria and the UK should be obliged to give the police their passwords. Home affairs ministers are currently conducting a centralisation of all European police databases.

Institutions and research programs of the European security architecture

In order to have more control over mass protest, for example during G8 summits, new

institutions and research programmes have been developed. European police forces conduct joint trainings and maneuvers to control demonstrations. In European police academies operational tactics for 'crowd management' are designed. The European Police Academy (CEPOL), based in Hampshire, UK plays a crucial role: 'CEPOL's mission is to bring together senior police officers from police forces across Europe – essentially to support the development of a network – and encourage cross-border cooperation in the fight against crime, public security and law and order, by organising training activities and research findings'.

Following the summit protests in Genoa and Gothenburg in 2001, in 2004 the EU initiated the research programme, 'Coordinating National Research Programmes on Security during Major Events in Europe' (EU-SEC). EU-SEC coordinates police departments of EU member states and Europol and publishes a handbook for summit protests. Police are advised to observe protest movements, to exchange data, to enforce travel bans, and to undertake aggressive media strategies in order to delegitimise resistance. In the form of questionnaires, information is gathered about European groups and individuals, their action forms, websites, mail addresses, international contacts, preferred travel routes, means of transport and accommodation.

EU-SEC is coordinated by the UN working group 'International Permanent Observatory on Security during Major Events' (IPO), based in the Italian city Turin. IPO advises governments on the appropriate security architecture for major events. IPO services are free. At the moment, IPO is putting together a 'Handbook for G8 states'. Official operational areas since its foundation in 2006 have so far been the G8 summits in St Petersburg and Heiligendamm, the World Bank/IMF summit in Singapore, and the Asia-pacific Economic Cooperation (APEC) meeting in Vietnam. Also, the Olympic Games 2008 in Beijing and the G8 summit 2008 in Japan were coached by IPO.

Border control: the militarisation of migration control

With the extension of the EU member states and the abolition of border controls, the new external EU borders are being technically upgraded. They include nightview technology, automatic analysis of video surveillance and high frequency cables that can measure and communicate the water concentration of nearby bodies. New joint headquarters have come into existence. Through the extension of the Schengen Information System II (SIS II), more data is available to police forces. Fingerprints and biometrical data of migrants are to be stored in the Visa Information System (VIS). Home affairs ministers complain about the insufficient police control of migrants and have demanded the use of RFID chips (chips with radio waves) in passports. These chips could, for example, acoustically identify the bearers of an expired visa, without this person actually having to show his/her passport.

With the creation of the 'Border Control Agency Frontex' in Warsaw, EU-wide 'migration control' now has another pillar.The General Director, Ilkha Laitinen, a Finnish border officer, summarises the 'Integrated Border Management' of Frontex in the following way, 'All those who don't deserve to be and whom one does not want to have on one's

2008-12-27 | 21:45 CET | 01:00 h | Saal 3 | lecture | Society | See paper on p. 261!

territory, have to be stopped.' In a 'risk analysis center' prognoses of waves of migration are undertaken, information is passed to the relevant border police departments and concrete measures are 'recommended'. Frontex has a 'central technical toolbox' for member states' control and surveillance of external borders. Frontex conducts operations together with national police forces ('Frontex Joint Support Teams'). Although Frontex has no forces of its own to fight migration, there has been an extensive increase in the arsenal of border forces of member states. The Italian Carabinieri for instance have new boats, helicopters and surveillance technology. According to its own publications, 115 boats, 27 helicopters, and 21 aeroplanes are documented in the central register of Frontex. Besides trainings, Frontex also undertakes research programs. For example, they research and recommend the use of 'micro-helicopters' for border observation. Director Laitinen has expressed his wish for Frontex to have more of its own resources and operative forces in the future.

Police and combatting counter-insurgency abroad

The Lissabon Treaty also addresses 'reforms' in the field of military affairs. The 'European Security and Defense Policy' asks for a 'gradual improvement of military capacities'. The Lissabon Treaty also plans 'reforms' within the field of military politics. The aim os for the EU to have armed units at its disposal by 2010. In January 2007 the first EU Battlegroup was declared fully operational; in 2006 such a unit was already considerably involved in the EU military deployment in Congo. There are also means for intervention in 'third states' that are much less visible: The 'European Gendarmerie Forces' (EGF). The EGF is a paramilitary police unit founded and developed at the G8 summits in 2002 and 2004. It should be able to mobilise 3 000 police officers within 4 weeks. Forces are so far provided by the Netherlands, France, Spain, Italy and Portugal. The EGF is supposed to take over police control after military deployments in crisis areas, as well as ensure 'public order' during the 'occurrence of public unrests'. The non-domestic deployment of police forces is considered a 'civilian instrument'. So far, maintaining 'public order' in 'third states' has been the task of the military, although it always has cooperated with police units. For example in Bosnia, members of the German Army were trained by Italian Carabinieri. The official tasks of the EGF include 'the entire spectrum of police deployments, civilian authority and military command, control of local police authorities, criminal investigation activities, activities for the provision of secret intelligence, property protection' etc. The statute of the EGF does not exclude a deployment within the EU. The headquarters of the EGF are located in the Italian city of Vicenza at a Carabinieri base. Likewise, in Vicenza the EGF have their own academy (COESPU) where their own forces as well as units of other countries are trained. The academy is financed by the G8 states. Also, senior police officers of Pakistan and Kenya have undergone COEPSU training in 'riot control'.

The significance for radical movements

'The distinction between international law in times of peace and in times of war is no longer appropriate in the face of new threats', Schäuble, the German Minister of Home Affairs has stated. The German chancellor and the head of the Federal Criminal

2008-12-27 | 21:45 CET | 01:00 h | Saal 3 | lecture | Society | See paper on p. 261!

Investigation police have further conceded that, 'the separation between internal and external security is obsolete'. What do these developments mean for the political practice of radical movements in general and for the alterglobalisation movement specifically, except 'even more repression'? A debate about repression should be an integral part of the practice of radical movements. It is clear that the margins for left interventions have not increased in light of and after 9/11. Nonetheless, we think that it is not only the speed and the degree of repressive measures that has changed. The entire social matrix within which radical left politics is situated is shifting. The quality of surveillance and social control has taken on another form. Apart from technological developments, above all this has to do with the transnational coordination of control agencies and the 'interdependency of internal and external security'.

But we see an opportunity in using this continued narrowing of the freedom of movement as a chance to build new alliances that will bring about broad social debates and unexpected interventions. A conjunction of classical antimilitarism, antirepression, and migration politics is a clear option. The degree to which these new measures and institutions touch upon the daily life of every European should offer sufficient starting points for a practice of proactive disobedience against this evolving European Security Architecture.

Against the European Security Architecture

The decision to mark the 60th anniversary of NATO with jointly hosted celebrations in Strasbourg (France) and Kehl (Germany) has already caused a great deal of activity amongst the antimilitarist left in a number of countries in Europe. The established peace movements in France and Germany plan to focus their protests on the Afghanistan war. However, the NATO summit would also be an excellent opportunity to draw attention to the complex structure of the 'global security architecture' with its participating institutions. Military and police forces currently maintain a repertoire of repressive instruments based on new technological developments. Computer-supported commandos, investigative software, warmth and body fluid detectors at national borders, tasers etc. Military and police remits are being ever more synchronised, both on a legislative basis and through joint operations, but also with the creation of common organisations such as the 'European Gendarmerie Force' based in Vicenza. 'Eurocorps, the French Foreign Legion and the central Schengen Information System are all located in Strasbourg, where next year's NATO conference is to be held. These facts provide ample

Gipfelsoli gipfelsoli@nadir.org
Gipfelsoli is participating in two websites, publishing a newsletter and help to spread information via workshops. See http://www.gipfelsoli.org and http://euro-police.noblogs.org.

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-27 | 23:00 CET | 01:00 h | Saal 1 | lecture | Hacking

Why were we so vulnerable to the DNS vulnerability?

SSL wasn't enough. Encryption is nonexistent. Autoupdaters are horribly broken. Why is all this the case?

> http://events.ccc.de/congress/2008/Fahrplan/events/2906.en.html

Effugas

2008-12-27 | 23:00 CET | 01:00 h | Saal 2 | lecture | Hacking

[redacted]

coreboot: Beyond The Final Frontier

Open source BIOS replacement with a radical approach to boot.

The BIOS and it's successor EFI are considered by many to be the final frontier for open source software in commodity PCs. This talk describes the BIOS replacement coreboot (formerly LinuxBIOS) and the projects surrounding it.

The closed nature of traditional firmware is starting to cause concern even on the government level, as awareness for BIOS malware risks is increasing.

The presentation describes coreboot, supplementary tools such as buildrom, flashrom, superiotool and nvramtool, and some popular payloads that combine with coreboot to make up the firmware: FILO, EtherBoot, SeaBIOS, Memtest86, tint, Linux, coreinfo, bayou and libpayload featuring tinycurses, which turns simple applications into instant-on appliances. Finally there will be a demonstration of coreboot running on hardware.

http://www.coreboot.org/Welcome_to_coreboot coreboot

> http://events.ccc.de/congress/2008/Fahrplan/events/2970.en.html

Peter Stuge peter@stuge.se

Peter is a self-employed consultant with many years of experience in software, hardware, security, networking, databases and more. An early hacker, he was taking electrical apparatus apart at the age of 4 and he actually put one or two back together too! He started programming at 9 and got his first soldering iron for his 10th birthday after collecting LEDs for a year.

2008-12-27 | 23:00 CET | 01:00 h | Saal 3 | lecture | Hacking

Messing Around with Garage Doors

Breaking Remote Keyless Entry Systems with Power Analysis

We demonstrate a complete break of the KeeLoq crypto-system. Thanks to Power Analysis, even non-specialists can gain access to objects secured by a KeeLoq access control system.

KeeLoq remote keyless entry (RKE) systems are widely used for access control purposes such as garage openers or car door systems. The talk will present the first successful differential power analysis (DPA) attacks on numerous commercially available products employing KeeLoq code hopping. They allow for efficiently revealing both the secret key of a remote transmitter and the manufacturer key stored in a receiver. As a result, a remote control can be cloned from only ten power traces, allowing for a practical key recovery in few minutes.
After extracting the manufacturer key once, with similar techniques, it is possible to recover the secret key of a remote control and replicate it from a distance, just by eavesdropping on at most two messages. This key-cloning without physical access to the device has serious real-world security implications, as the technically challenging part can be outsourced to specialists. During the talk, the attack will be practically performed. Finally, it will be shown how to take over control of a KeeLoq access control system, i. e., lock out a legitimate user while the attacker may still open the door.

> http://events.ccc.de/congress/2008/Fahrplan/events/3030.en.html

Timo Kasper tkasper@crypto.rub.de
Timo Kasper is a researcher in the field of embedded security and member of the Embedded Security (EMSEC) group at the Ruhr-University Bochum, Germany.

Thomas Eisenbarth eisenbarth@crypto.rub.de
Thomas Eisenbarth is a Ph.D. candidate at the Department of Electrical Engineering at Ruhr University Bochum.

2008-12-27 | 24:00 CET | 01:00 h | Saal 1 | lecture | Culture

Kurt Gödel – I do not fit into this century

Ein audiovisuelles Live-Feature

Manche bezeichnen ihn als größten Logiker seit Aristoteles: Der 1906 geborene Wiener Mathematiker Kurt Gödel rührte ab 1930 mit seinen Unvollständigkeitssätzen an den Grundfesten der Mathematik. Er wies nach, daß es in jedem formalen logischen System Fragen gibt, die unentscheidbar sind. Sein Arbeitsleben verbrachte der Wissenschaftler, der wie viele Kollegen aus Europa fliehen mußte, am berühmten Institute for Advanced Study in Princeton – dem Mekka der modernen Mathematik. Der introvertierte Mensch Kurt Gödel schwankte dabei Zeit seines Lebens zwischen Genie und Wahnsinn, hatte zahlreiche Neurosen und eine ausgeprägte Paranoia.

Das live gesprochene audiovisuelle Feature der Hörspielwerkstatt der Humboldt-Universität zu Berlin dokumentiert Leben und Werk des Mathematikers in Wort und Bild mit Musik.

http://46halbe.org/audio.html Texte früherer Features

> http://events.ccc.de/congress/2008/Fahrplan/events/3050.en.html

46halbe

Marcus Richter

Ina Kwasniewski

Kai Kittler

Day 2

2008-12-28 | 11:30 CET | 02:15 h | Saal 1 | lecture | Community

Jahresrückblick

Die Themen des CCC im Jahr 2008

Es war mal wieder ein bewegtes Jahr für den CCC. Was alles passiert ist, werden wir in der gebotenen Kürze berichten.

Der CCC konnte im Geschäftsjahr 2008 in allen Sparten positive Ergebnisse ausweisen. Neben den gesetzlichen Zumutungen konnten die Entwicklungen moderner Technik weiter verstärkt durchdrungen und kommentiert werden. Damit ist sichergestellt, dass der CCC für seine Mitglieder und die interessierte Öffentlichkeit auch in Zukunft ein solider und verläßlicher Partner bleibt.

Der Bekanntheitsgrad stieg um 21% auf nunmehr 81%, während der Beliebtheitsgrad des CCC bei den politischen Entscheidungsträgern um 18% auf nun 29% sank. Dies entspricht in etwa den prozentualen Ergebnissen der Oppositionsparteien bei Wahlen. Die Rückstellungen noch nicht veröffentlichter Hacks wurden dabei weiter erhöht.

Durch unsere seit Jahren auf Stabilität und Unabhängigkeit ausgelegte Außenkommunikation war es dem CCC möglich, die durchschnittliche Schlagkraft betriebener Kampagnen von 6% über den Branchenschnitt von ca. 5% zu steigern. Der Verpeilfaktor konnte von 79,8% auf 79,7% gesenkt werden – ein weiterer Zehntelprozent-Schritt in Richtung einer strahlenden Zukunft!

Hervorragender Service – zufriedene Mitglieder

In von unabhängigen Instituten durchgeführten Umfragen hat der CCC auch 2008 vor allem in den Bereichen Beratung und Service für Politiker wieder Bestnoten erhalten. Diese Fachberatung sowie die Kundennähe der Außenstellen des CCC haben ihre Wirkung gezeigt: Gegenüber dem Vorjahr konnten rund 22% mehr Politiker bespaßt und informiert werden. Da gemessen am Vorjahr ebenfalls wieder mehr Veranstaltungen für die Mitglieder angeboten werden konnten, betrug der Unzufriedenheitsgrad der Mitglieder lediglich 0,9%. Über alle Sparten hinweg konnte der Mitgliederbestand weiter erhöht werden.

Der CCC konzentrierte sich weiterhin darauf, tatsächliche Informationen an die interessierte Öffentlichkeit weiterzugeben, anstatt die Entsolidarisierung und weitere Überwachung zu fördern, wie das bei der Lancierung vieler neuer Gesetze ursprünglich der Fall gewesen wäre.

Laut einer Umfrage vom Sommer 2008 schätzten die Befragten die Durchschlagskraft von Argumenten auf horrende 22%, eine Zahl, die weitaus höher als der Schnitt bei Politikerreden von 5,4% liegt. In dieser Hinsicht ist dennoch noch viel Aufklärungsarbeit nötig.

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-28 | 11:30 CET | 02:15 h | Saal 1 | lecture | Community

> http://events.ccc.de/congress/2008/Fahrplan/events/3024.en.html

46halbe

andy

Frank Rieger

frankro

Erdgeist

2008-12-28 | 11:30 CET | 01:00 h | Saal 2 | lightning | Community

Lightning Talks 1

5 mins of fame

5 minutes for every speaker. Learn about the good, the bad, and the ugly - in software, hardware, projects, and more.

Give a lightning fast talk about your favourite project, program, system - and thereby find people with the same interest to proceed and promote it. Alternatively - give us a good rant about something and give us some good reasons why it should die. ;)

Get right at it, don't waste time by explaining too much, get the main points across, and then let us know how to contact you on the congress for a talk!

Whatever you do - please practise it, and don't be boring. Or else. You have been warned! :-P

> http://events.ccc.de/congress/2008/Fahrplan/events/3048.en.html

SvenG

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-28 | 11:30 CET | 01:00 h | Saal 3 | lecture | Society

Embracing Post-Privacy

Optimism towards a future where there is 'Nothing to hide'

The breaking away of privacy in the digital world is often understood as something dangerous, and for good reasons. But could there be opportunities in it, too? Do the current cultural and technological trends only dissolve the protected area of privacy, or could they dissolve as well the pressures that privacy is supposed to liberate us from? What if we witness a transformation of civilization so profound that terms like 'private' and 'public' lose their meaning altogether? Maybe we won't need 'privacy' at all in the future because we will value other, new liberties more strongly?

In the digital world, more and more data is accumulated about us. More and more methods of datamining are invented to extract information from these data. The youth grows up enjoying informational exhibitionism to a degree many find irresponsible. Ever greater parts of life are integrated into the global public information stream. Will privacy end? If so, what about liberty? We have to look closely at the value of privacy. What does it do for values like freedom, individualism or intimacy? Why is this protected area of privacy necessary?

The conditions of privacy are rapidly changing. We have to evaluate these changes with a perspective that does justice to new modes of identity, sociality and culture: Why hide your personal weirdnesses if 21st century society thrives on difference and originality instead of conformism and predictability? What identity is there to keep private if 'identity' is more and more what you externalize from yourself into the internet? Is privacy worth missing out on participation in the global 'hive mind' and the 'ambient intimacy' of every mind connected with every other mind?

Such questions may sound utopian and/or crazy. They may sound irresponsible, considering anti-privacy trends that may seem much more real and dangerous -- like the surveillance state. But even if you disagree with their validity, they may provoke deeper thinking about the state and value of privacy in a world that is changing more and more rapidly -- and that could hardly be a bad thing.

> http://events.ccc.de/congress/2008/Fahrplan/events/2979.en.html

Christian Heller / plomlompom

Christian Heller / 'plomlompom' is a 23-years-old futurist, blogger and film critic from Berlin, Germany.

2008-12-28 | 12:45 CET | 01:00 h | Saal 2 | lecture | Science | See paper on p. 243!

Climate Change - State of the Science

We are in the midst of a major global warming, as witnessed not just by temperature measurements, but also for example by the record loss of Arctic sea ice in 2007 and 2008. This year, both the Northwest Passage and the Northeast Passage in the Arctic were open for ships to pass through for the first time in living memory.

What are the causes of this warming? And how will it affect sea level, tropical storms and other aspects of the climate system? And can we stop this warming, and how? These topics will be discussed based on the most recent scientific results, by one of the leading climate scientists in the world.

http://www.pik-potsdam.de/~stefan/	*home page of speaker*
http://www.realclimate.org/	*RealClimate Weblog*
http://www.wissenslogs.de/wblogs/blog/klimalounge	*KlimaLounge Weblog (in German)*

> http://events.ccc.de/congress/2008/Fahrplan/events/2863.en.html

Rahmstorf
Stefan Rahmstorf is one of the authors of the IPCC climate report published last year, as well as a member of the Advisory Council on Global Change of the German government. He has published over fifty peer-reviewed publications in climate science, as well as four books on the topic (two of which are still in press at the moment).

2008-12-28 | 12:45 CET | 01:00 h | Saal 3 | lecture | Hacking | See paper on p. 171!

Full-Disk-Encryption Crash-Course

Everything to hide

This is not a hacking presentation, no vulnerabilities are presented. It's a crash-course in full-disk-encryption ('FDE') concepts, products and implementation aspects. An overview of both commercial and open-source offerings for Windows, Linux, and MacOSX is given. A (programmer's) look at the open-source solutions concludes the presentation.

Full-Disk-Encryption is an important aspect of data security and everyone should use an appropriate solution to protect their (especially mobile) systems and data. This lecture covers the technology behind Full-Disk-Encryption software products.

The established technical architectures of software solutions for Microsoft Windows and Linux are presented in this lecture: Pre-Boot-Authentication, Encryption driver, In-place filesystem encryption.

An overview of commercial products and open-source offerings for Windows, Linux and OSX is given. Distinguishing features of specific products and additional topics are covered, including: TPM support (OS binding and key storage), Multi-disk support, Threats.

The last segment of the lecture focuses on open-source solutions: TrueCrypt's volume specifications, TrueCrypt's hidden volume capabilities, Comparison of in-place filesystem encryption implementations of TrueCrypt and DiskCryptor. A feature wish-list for open-source Full-Disk-Encryption solutions is presented completes the lecture.

http://blog.akkaya.de/jpabel Juergen Pabel's Blog

> http://events.ccc.de/congress/2008/Fahrplan/events/2882.en.html

Juergen Pabel
Juergen is working as an IT-Security Consultant and likes to play Rugby.

2008-12-28 | 14:00 CET | 01:00 h | Saal 1 | lecture | Hacking

Attacking Rich Internet Applications

Not your mother's XSS bugs

This presentation will examine the largely underresearched topic of rich internet applications (RIAs) security in the hopes of illustrating how the complex interactions with their executing environment, and general bad security practices, can lead to exploitable applications.

In recent years rich internet applications (RIAs) have become the mainstay of large internet applications and are becoming increasingly attractive to the industry due to their similarity to desktop applications. Furthermore their user of exsting web technologies such as HTTP, HTML/XML and Javascript/Actionscript make them attractive options to companies with existing web developers.

Unfortunately the use of existing technologies brings with it the burden of existing ways to write vulnerable code, but adds yet more ways. This presentation will examine the largely underresearched topic of RIA security in the hopes of illustrating how the complex interactions with their executing environment, and general bad security practices, can lead to exploitable applications.

> http://events.ccc.de/congress/2008/Fahrplan/events/2893.en.html

kuza55
kuza55 Hacks things, mostly new shiny web things; he also likes telling people about it.

Stefano Di Paola

2008-12-28 | 14:00 CET | 01:00 h | Saal 2 | lecture | Science

Hacking the Atmosphere

How to exploit the weather and fly for free

Birds, glider pilots, and recently UAVs can exploit a variety of weather effects in order to gain altitude, remain airborne and travel long distances all with no power input – effectively, hacking the atmosphere to fly for free. This talk will explain the aircraft, techniques, meteorology, hardware and software that we use to achieve this. In the process I will show why the sport of gliding may be of interest to hackers, and explain how you too can get involved in this highly rewarding and low-cost form of flying.

Introduction – a brief history and philosophical motivation.

Vulnerabilities – an introduction to the different mechanisms by which we can extract energy from the atmosphere including ridge, thermal and wave lift and dynamic soaring.

Hardware – comparison of the various types of soaring aircraft including sailplanes, hang gliders and paragliders, R/C gliders and recent autonomous soaring UAV projects, along with instruments and other equipment.

Software – a look at the increasing array of both ground-based and in-flight software used to aid soaring, much of it now in the form of open source projects developed by and for pilots.

Results – an overview of modern gliding and its accomplishments, and some tales from my own experience as a pilot and instructor.

> http://events.ccc.de/congress/2008/Fahrplan/events/2940.en.html

Martin Ling martin-ccc@earth.li
Electronic engineer and glider pilot.

2008-12-28 | 14:00 CET | 01:00 h | Saal 3 | lecture | Hacking

Exploiting Symbian

Symbian Exploit and Shellcode Development

SymbianOS is one of the major smart phone operating system and has been around for many years still exploitation has not been researched yet. The lack of proper exploitation techniques is mostly due to the fact that until the recent introduction of PIPS/OpenC (a POSIX API port) SymbianOS did not have the means for programmers to EASILY write insecure code.

The presented work will show that now it is possible to exploit buffer overflows on Symbian like on any other (mobile) platform. To do this we will show some proof-of-concept exploits and provide an overview on writing shellcode for SymbianOS.

Further we will show some short comings of the Symbian security model and discuss the possible impact. To do this we will show that it is possible to create a piece of (mobile) malware that is capable to sign itself.

We believe vulnerability exploitation will become the next big issue on SymbianOS because the current version of Symbian only permits installation of signed applications thereby shutting out currently existing Symbian worms. We believe worm authors will adapt soon.

http://www.mulliner.org/symbian/

> http://events.ccc.de/congress/2008/Fahrplan/events/2832.en.html

Collin Mulliner collin-25c3@mulliner.org
Security Researcher at day, Hacker at night. Loves gadgets, especially small ones without wires.

2008-12-28 | 16:00 CET | 01:00 h | Saal 1 | lecture | Hacking

Vulnerability discovery in encrypted closed source PHP applications

Security audits of PHP applications are usually performed on a source code basis. However sometimes vendors protect their source code by encrypting their applications with runtime (bytecode-)encryptors. When these tools are used source code analysis is no longer possible and because these tools change how PHP works internally, several greybox security scanning/fuzzing techniques relying on hooks fail, too.

This talk will show how different PHP (bytecode-)encryptions work, how the original bytecode can be recovered, how vulnerability discovery can still be performed with only the bytecode available and how feasible PHP bytecode decompilation is.

> http://events.ccc.de/congress/2008/Fahrplan/events/2678.en.html

ionic

2008-12-28 | 16:00 CET | 01:00 h | Saal 2 | lecture | Making | See paper on p. 219!

Algorithmic Music in a Box

Doing music with microcontrollers

Small devices like microcontrollers, coupled to a few buttons, knobs, encoders and LEDs, allow for a host of interesting and creative musical applications. Solder a few bits together, program a few lines, and you can build a deep device to support your musical exploration. This lecture will show you quickly how the hardware and code works, and then focus on a few interesting applications: controllers, sequencers, sound generators. The workshop will allow you to build your own crazy ideas.

If you look at 'real' instruments, their principle is very simple most of the time: press a key, get a sound. Pluck a string, get a sound. This simple principle is also what makes their complexity. Every aspect of the interaction with these instruments has been explored. Modern music making software on the other hand often displays thousands of different features, which in the end often works against the musician by blurring his focus.

Some of the most sought after electronic devices are from the 80ies, and feature pretty simple (by today's standards) controlling units and interfaces. Instead of focusing on the 'next big thing', like touchscreens, new sensors and technologies, building musical instruments with limited resources allows the musician to go 'deep', to think about interaction in a very conscious way. Devices like the monome or the tenori on, featuring just buttons and leds, are following this direction.

Devices like these are pretty easy to build at home, and then can be built robustly enough to find their way into live performances and daily studio use. This event's lecture is about the overall design of devices I built, and shows how the different hardware elements work (microcontroller, buttons, knobs, displays), and then focuses on the flexible ways you can use these simple elements to build interesting instruments. The workshop will then allow participants to create their own instrument based on a simple hardware platform and to take it home with them.

http://ruinwesen.com/blog *wesen blog*
http://ruinwesen.com *wesens hardware*

> http://events.ccc.de/congress/2008/Fahrplan/events/2843.en.html

wesen wesen@ruinwesen.com

2008-12-28 | 16:00 CET | 01:00 h | Saal 3 | lecture | Hacking

Handschellen hacken

Essentielles Grundwissen für alle, die nichts zu verbergen hatten

Jeder kann auf Youtube ansehen, wie man normale Handschellen mit einer Büroklammer öffnet. Aber es gibt verschiedenste Hochsicherheitsmodelle mit deutlich komplizierteren Schlössern, die nur darauf warten, vom Sperrsport entdeckt zu werden...

Den meisten ist bekannt, dass normale Polizeihandschellen keine echte Herausforderung für den ambitionierten Lockpicker sind, obwohl es natürlich hilft, die kleinen Unterschiede der verschiedenen Modelle zu kennen. Wesentlich interessanter wird es jedoch, wenn man sich mit den verschiedenen Hochsicherheitsmodellen beschäftigt, die vor allem für Risikotransporte eingesetzt werden, aber in manchen Gegenden auch bei normalen Streifenpolizisten anzutreffen sind.

Dieser Vortrag verschafft einen umfassenden Überblick über die Vielzahl verschiedener in Handschellen eingesetzten Schließmechanismen – und ihre Schwächen. Dabei geht es nicht nur um das Picken diverser Chubbschlösser und Stiftzylinder, die in dem Umfeld durchaus zum Einsatz kommen, sondern auch um überraschende Umgehungstechniken, die das mehr oder weniger raffinierte Schließsystem links liegen lassen. Neben dem inzwischen 'klassischen' Angriff auf die amerikanischen Handschellen mit Medeco-Zylinder, der bereits auf der HOPE in New York viele Freunde fand, werden auch einige neue oder wenig bekannte Angriffe auf in Europa verbreitete Modelle vorgeführt und erklärt.

Nach dem Vortrag bietet sich im Lockpickingbereich die Gelegenheit, das theoretisch erlernte an einer Vielzahl mitgebrachter Sportgeräte 'Hands on' auszuprobieren.

> http://events.ccc.de/congress/2008/Fahrplan/events/2963.en.html

Ray

Ray's been around for years. Besides lockpicking and random number generating he's interested in linux, electronics and stuff.

2008-12-28 | 17:15 CET | 01:00 h | Saal 1 | lecture | Hacking

Anatomy of smartphone hardware

Dissecting contemporary cellphone hardware

Do you know the architecture of contemporary mobile phone hardware? This presentation will explain about the individual major building blocks and overall architecture of contemporary GSM and UMTS smartphones.

We will start from a general block diagram level and then look at actual chipsets used in mobile devices, ranging from SoC to RAM and flash memory technologies, Bluetooth, Mobile WiFi chipsets, busses/protocols as well as the GSM baseband side.

The main focus will be about the OpenMoko Freerunner (GTA02) hardware, since the schematics are open and can be used for reference during the lecture. However, we will also look into tighter integrated components of various vendors like Qualcomms MSM7xxx, Samsung S3C64xx, TI OMAP35xx and others.

> http://events.ccc.de/congress/2008/Fahrplan/events/3008.en.html

Harald Welte laforge@gnumonks.org

2008-12-28 | 17:15 CET | 01:00 h | Saal 2 | lecture | Hacking | See paper on p. 183!

Security of MICA-based wireless sensor networks*

Firstly, we mention an initial qualitative risk assessment, carried out by interviewing the operating manager of a large suspension bridge and a contractor responsible for part of a large subway tunnel network who want to use wireless sensor networks. The core of the talk deals with assessing the practical security of the particular COTS system adopted by our team, the Crossbow MICAz motes running TinyOS or XMesh, together with the Stargate gateway: we designed and implemented a variety of attacks on this system and we discuss the security problems we found, together with appropriate fixes where possible. While some of our attacks exploit generally known vulnerabilities, others like selective jamming and power exhaustion through routing table manipulation are original and interesting in their own right. In section we also demonstrate how an attacker can undetectably alter messages in an IEEE 802.15.4 radio environment.

It is available in our paper 'Steel, Cast Iron and Concrete: Security Engineering for Real World Wireless Sensor Networks' published in ACNS 2008 (Applied Cryptography and Network Security) conference.

> http://events.ccc.de/congress/2008/Fahrplan/events/2831.en.html

Dan Cvrcek dancvrcek@gmail.com

2008-12-28 | 17:15 CET | 01:00 h | Saal 3 | lecture | Hacking

TCP Denial of Service Vulnerabilities

Accepting the Partial Disclosure Challenge

The Transmission Control Protocol (TCP) is one of the fundamental protocols used in today's communication networks. Recently, there has been an increased discussion on possible Denial of Service attacks against TCP-based services, which has largely been triggered by the partial disclosure of several vulnerabilities by the security company Outpost24. This talk will present several TCP vulnerabilities in an attempt to find out just what they found.

This year, vulnerabilities have been identified in the specifications of various core network protocols. This included BGP, DNS and TCP. Accompanying these wide-ranging discoveries, a new form of vulnerability disclosure named 'partial disclosure' has been introduced. In practice, this means that the public knows that there is something wrong, yet, it is uninformed about the details. This, of course, can be understood as a challenge to find out just what could be wrong, which is what we at Recurity Labs did after the Denial of Service vulnerabilities in TCP had been announced.

This talk will present known vulnerabilities in the protocol, which have been receiving rather sparse media-attention, as well as some attacks we have been working on during our research. Additionally, we hope to provide sufficient background information on the protocol's fundamental weaknesses to motivate further research on the subject. We argue that certain assumptions made by the protocol engineers almost 30 years ago do not hold in today's networks and that most possible Denial of Service attacks against TCP can be derived from these assumptions.

> http://events.ccc.de/congress/2008/Fahrplan/events/2909.en.html

Fabian Yamaguchi

2008-12-28 | 18:30 CET | 01:00 h | Saal 1 | lecture | Hacking

Short Attention Span Security

A little of everything

Working as a security consultant means that you get to see everyone's dirty laundry. However, it also means a hectic schedule and restrictive confidentiality agreements. Without violating my NDA, here's a set of turbo-talks looking at some new tricks for some new technologies and a look at some lucrative new attack surfaces that will become much more prevalent in the coming year. Topics will include: Script Injection in Flex, EFI Rootkits, static analysis with Dehydra, and pattern-matching hex editors.

Things I want to talk about (details below):
• EFI Rootkits
• Bypassing MS anti-XSS libraries
• Script injection in Flex
• Pattern-matching hex editors
• Static analysis with Dehydra
• Auto-WEP key cracking with ITX
• Porting Network Security Tools to the iPhone

Along with this, I can make some code available for the hex editor, a bunch of iPhone security apps as an Installer repository, some Dehydra stuff and the source for my little WEP-cracking ITX box.

I want to strip out all the usual introduction and fluff and do 5-7 turbo talks (with two of them being extremely short). Or one of these could be done as a separate turbo talk.

EFI Rootkits
In the next year, every major chip manufacturer will ship boards that use EFI. This brings new life to the old idea of PCI Option ROM rootkits, which can now easily access libraries that provide filesystem access as well as a full network stack. What features of EFI make this easy? What are the constraints on an EFI rootkit? How could this be mitigated as an attack vector?

Bypassing MS anti-XSS libraries
This is a quick one. There is a bug in the Microsoft implementation of libxml, such that the attributes of start and end tags are merged. This means that Internet Explorer respects XML attributes on end tags. There is a particular Microsoft anti-XSS library which looks for an ' <' followed by any letter. It allows a ' <' followed by a '/' however. To bypass this library, simply put your script in an end tag attribute, like so:
</a style='background:expression(alert(document.cookie))'>

Script injection in Flex
Since the provided user controls handle input encoding, injections are scarcer, but still available. One less conventional method I found relies on a bug in Internet Explorer. On a

2008-12-28 | 18:30 CET | 01:00 h | Saal 1 | lecture | Hacking

web application that allows file uploads, perhaps attachments, you can upload an HTML file containing the injection script.

When this attachment is viewed in Firefox, it will behave correctly and download the file first and then view it in a local file script context. In IE however, the downloaded HTML file is viewed with the script context of the site from which it was downloaded!

So once you have a script injection, Flex can make life difficult with URL scrambling - kind of like ASLR for web apps. Your injected script has to make several requests via AJAX to retrieve and parse the URL mapping for the current session. I have an example script.

Static analysis with Dehydra
A new patch for GCC from Mozilla, Dehydra, allows the scripting of custom static analysis rules using Javascript via the SpiderMonkey engine. How does this make your life easier on the first two days of a code audit? Interesting semantic searches to perform on C++ code bases, advantages and limitations of this approach.

Pattern-matching hex editors
Introducing my toy pattern-matching hex editor, haxedit, which can visually demonstrate the effectiveness of various pattern-matching algorithms on arbitrary binaries.

Auto-WEP key cracking with ITX
This has become so trivial, people are playing for time with average scores under 3 minutes. Tips and tricks for working around the idiosyncrasies of airtools in an embedded environment.

Porting Network Security Tools to the iPhone
Probably drop this, since it's all on the App Store now ...

> http://events.ccc.de/congress/2008/Fahrplan/events/2734.en.html

Ben Kurtz
Ben Kurtz has spent the last year in the trenches as a security consultant in Seattle. In previous incarnations, he worked on avionics, power plants, visualizations, automated network death machines, and once knitted a quite nice scarf.

Ben spoke at Defcons 13 and 15, and at CCCamp07.

2008-12-28 | 18:30 CET | 01:00 h | Saal 2 | lecture | Society | See paper on p. 309!

The Trust Situation

Why the idea of data protection slowly turns out to be defective

In many social situations, people start to adjust their behaviour due to surveillance. Inspired by more and more cases of breaches of data protection regulations, an erosion of trust into these regulations and those who forfeit them can be seen. The consequences of this are grim. Either we abolish surveillance technologies or the idea of 'informational self-determination'.

Surveillance is beginning to show us some first substantial side effects. As its mere technological existence is sufficient to evoke the impression of potential identification in any situation, many people in special situations who fear repercussions emerging from such identification are beginning to be substantially manipulated by surveillance. People in need of aid such as troubled families or drug addicts stop seeking aid as they fear that they will be identified and observed closely henceforth. Informants of the press cannot rely on their anonymity anymore as they know that this can just not be guaranteed anymore. The same applies to witnesses and defendants. They fear telling details of their cases to their own lawyers as they know for a fact that those can be wiretapped too.

Thus what we see is that many social arrangements needed in a just and democratic society or arranged in solidarity actually start to crumble. A decisive thing to note about this now is that data protection regulations do not prevent these people from alterting their behaviour anymore. As data protection needs a situation of trust into many things - the law, science, technology, companies - and as this trust is betrayed more and more often, the mere technological possibility of surveillance becomes more and more sufficient to produce these effects. Thus the idea of data protection to ex post facto regulate the use of surveillance technologies and data is now at its limits. Admitting this, the consequences are grim. Either the technologies themselves have to be abolished again - and that's not going to happen - or the idea of 'informational self-determination' has to be given up. And that's 'Goodbye freedom'.

> http://events.ccc.de/congress/2008/Fahrplan/events/2665.en.html

Sandro Gaycken

2008-12-28 | 18:30 CET | 01:00 h | Saal 3 | lecture | Making

Scalable Swarm Robotics

Formica: a cheap, open research platform

The topic of swarm robotics will be introduced, including the current state of the art and some current research platforms. The problems of scalability in robot swarms will be discussed, particularly of programming and maintaining a large group of robots. The Formica platform represents a novel, very low cost approach to swarm robotics. Its design and implementation will be described, and the lecture will culminate in a live demonstration of a swarm of 25 robots cooperating on a task.

Swarm robotics is a hot research area. In cases such as earthquake rescue or planetary exploration, a swarm of cheap, simple robots may benefit from redundancy and distributed problem-solving. However, the cost of current robotics platforms prohibits experimentation with swarms numbering more than a few tens of units. As a result, the practicalities of software and hardware maintenance in large swarms are yet to be addressed.

At the University of Southampton, four colleagues and I developed a small, low-cost platform for swarm robotics research. We named it Formica. 25 robots 25x25x15mm in size were designed and built, costing only £25 each. They are capable of infrared communication, sensing and reprogramming, autonomous charging, and can drive for around 2 hours before a recharge. We presented them at ALIFE XI, the 11th international conference on artificial life, where they attracted a lot of attention from other researchers and the press. As a result, we have released the hardware and software under open source licenses to encourage further development in the community.

I would like to describe the design of the Formica platform and introduce some of the research it is used for. The swarm of 25 robots was assembled and tested by a small team of students in a single day. I will briefly describe how we overcame the problems of mass production on a very low budget. I will take questions from the audience, and encourage 25C3 attendees to hack on the platform. Finally, I will give a live demonstration of the swarm in action.

http://www.youtube.com/watch?v=65OZYWoYrtk A rough, early video of the swarm
http://news.bbc.co.uk/1/hi/technology/7549059.stm
BBC coverage of Formica at the ALIFE conference
http://www.electronicsweekly.com/Articles/2008/08/06/44275/southampton-university-designs-robots-for-mass-production.htm
Some more detailed press coverage
http://warrantyvoidifremoved.com/formica
A detailed overview of the project

> http://events.ccc.de/congress/2008/Fahrplan/events/2890.en.html

[redacted] nothing [redacted] to [redacted] hide [redacted]

2008-12-28 | 18:30 CET | 01:00 h | Saal 3 | lecture | Making

Jeff Gough jeff@warrantyvoidifremoved.com

I am a PhD student at the University of Southampton, UK. I am researching novel substrates for computation, and biologically inspired robotics. I am an advocate of open hardware and software, and a supporter of the hacker movement. I am becoming increasingly interested in the fusion of electronics, mechanics and art, and am seeking new projects which combine these elements. I enjoy metalwork, particularly precision machining, and make fine jewelry.

2008-12-28 | 20:30 CET | 01:00 h | Saal 1 | lecture | Making

Rapid Prototype Your Life

The time is now to make anything you can imagine

The tools are at hand to free you from the bonds of consumer slavery. No longer must you rely on distant and faceless factories or bow down before the false idols of mass produced consumer manufactured items. Never again look into the aisles of oblivion filled with mass produced products. Take rapid prototype manufacturing into your life and return to a time before corporations robbed you of our individualism. A cottage industry paradise awaits those with the digital skills and the means to acquire or build the machines that can actualize the items that exist now only in your imagination.

Presenting research into the potential to replace all consumer products with personally designed and built items utilizing rapid prototyping manufacturing.

This talk will present an overview of the rapid prototype machines available and take the audience on a visual adventure into the beautiful world of rapid prototyping machinery with enough luscious graphic detail to make even the most die-hard luddite salivate with lust for the dream fulfilling technology.

Projects to get beginners started and initiated into the rapid prototype lifestyle and familiar with the principals and techniques of rapid prototyping will be shared and source code provided online. These projects will be shared to get you started with rapid prototype machinery and break the cycle of consumer desperation. Many pathways will be presented leading to a future of wonderful digitally sharable objects.

> http://events.ccc.de/congress/2008/Fahrplan/events/3015.en.html

Bre

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-28 | 20:30 CET | 01:00 h | Saal 2 | lecture | Science

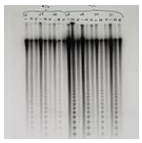

All your base(s) are belong to us

Dawn of the high-throughput DNA sequencing era

New DNA genotyping and sequencing technologies have recently advanced the possibilities for both mass and individual genomics by several orders of magnitude. The personal genome on DVD, genetic analysis of entire populations, and government DNA databases are but a few of the results of this process. The field is still accelerating, and the related computational challenges are enormous.

In the year 2000, completion of sequencing of the human genome was announced, a work taking decades, costing millions and involving hundreds of scientists around the world. Subsequent advances in DNA sequencing technologies have propelled the possibilities in the field to scales unthinkable a mere decade ago. The price of sequencing an entire human genome quickly approaches $1.000, and can be done by a few individuals and a single machine in a few days. Despite this, more powerful sequencing technologies are under development, and could simplify the process even further within the coming years.

Genotyping is a technology to quickly and cheaply analyze a DNA sample for potential SNPs (single nucleotide polymorphisms, aka point mutations) on a single plate (chip). Today's DNA chips can check for one million SNPs in a cheap and automated fashion. This allows to compare groups of thousands of people for specific markers. Applications for this technology range from finding resistance genes over evolutionary relations to the separation of an individual's DNA from a mixture of thousands of people.

Both technologies require new approaches in computational approaches and storage technology. Analysis is performed on massive computer clusters with thousands of CPUs. Data storage requirements are measured in petabytes, pushing hard disk storage to the limit.

In my talk, I will describe how we got here, how we handle the technological challenges involved, and what the future might hold.

http://www.sanger.ac.uk/ Where I work

> http://events.ccc.de/congress/2008/Fahrplan/events/3044.en.html

Magnus Manske
Long-time programmer (original author of MediaWiki), PhD in Biochemistry, working in Bioinformatics on large-scale DNA sequencing

2008-12-28 | 20:30 CET | 01:00 h | Saal 3 | lecture | Hacking

Banking Malware 101

Overview of Current Keylogger Threats

In the recent years, we observed a growing sophistication how credentials are stolen from compromised machines: the attackers use sophisticated keyloggers to control the victim's machine and use different techniques to steal the actual credentials. In this talk, we present an overview of this threat and empirical measurement results.

Nowadays, attackers often steal sensitive information from a victim's machine with the help of a keylogger that sends the stolen information to a so called dropzone. A dropzone is a publicly writable directory on a server in the Internet that serves as an exchange point for keylogger data: the malware running on a compromised machine sends all stolen credentials to the dropzone, where the attacker can pick them up and start to abuse them. Such an approach is more promising than 'traditional' phishing sites since the attacker can steal many more credentials from a single victim. In this talk, we present the results of an empirical study of this phenomenon, giving many details about these attacks we observed during the recent months.

In the first part of the talk, we provide a detailed overview of some of the most common keyloggers found in the wild. We focus on the two malware families ZeuS/Zbot and Limbo/Nethell and show how they propagate, what features they have, and how the actual dropzone works. Several other malware families will be briefly covered to cover a larger number of threats. Afterwards, we present several statistics and qualitative information for the keylogger data we found on some dropzones.

> http://events.ccc.de/congress/2008/Fahrplan/events/3020.en.html

tho thorsten.holz@gmail.com
Still working on his Ph.D. thesis...

2008-12-28 | 21:45 CET | 01:00 h | Saal 1 | lecture | Society

The Infinite Library

Storage and Access of Pornographic Information

Decades ago, Jorge Luis Borges wrote about infinite libraries and perfect memory with the slightly sad air of someone who'd seen those things and knew their faults. Today we work toward infinite libraries and perfect memory with little heed for the possible consequences. How could it be bad to have everything possible stored? To remember everything? I don't know that it will be bad, but I do know that it will be different from our current lives of loss and forgetting. Right now, storing pornography causes problems even for people who have nothing especially perverted to hide: A collection of pornography gets to the heart of what it means to be a private individual. As we move from mass media to individually produced media, from edited collections of porn (magazines, commercially produced films) to individual snapshots and youtube clips and stored bittorrents, the particularity of a collection of porn will be testimony to its owner's private set of tastes.

Of course, it has always been a pain to store pornography -- and so we have the cultural trope of a stash of magazines 'under the mattress' or in a box hidden in the closet. But as the sex industry shifts toward digital publication at every level, we might imagine that mere storage will become a problem of the past, or, at least, a problem related to legacy materials (books, magazines, videos, comic books, photographs, etc.). Cheap, massive storage media means no more problem, right?

Well, reviewers of porn find that they quickly amass more material than they will ever have time to peruse; librarians who need to provide access to controversial and poorly cataloged material end up overwhelmed; even casual collectors of pornography still need some way to keep track of what they have.

Toward that end, I am doing preliminary research on how people store and access their digital pornography collections. In my early interviews, I have already encountered a fascinating mix of responses; one person has said they store their porn 'in the cloud,' while another explained his detailed system for hiding digital porn files from his partner.

As I close, I will spend some time considering how we will store the pornography that isn't even being created yet. If science fiction author Charlie Stross is right, before long we will all be'life-logging' -- recording everything that happens to us, which of course would include all our sexual experiences. I think we might also one day be able to indulge in fully immersive AI-driven pornographic experiences (such as texting back-and-forth with artificially intelligent SMS-bots, sending texts and photos and audio to a perfectly responding far-away 'partner'), and we'll also want some way to keep those experiences.

We'll have everything stored -- but what will the social consequences be? If it is trivially easy to amass a porn stash so large that it cannot be 'consumed' in one person's lifetime, should a person with a large collection of pornography be considered a pervert? (Hint: I

2008-12-28 | 21:45 CET | 01:00 h | Saal 1 | lecture | Society

don't think so.) If everyone has so much porn, perhaps we have nothing to hide!

http://www.flickr.com/photos/lydiat/309105728/event image attribution

> http://events.ccc.de/congress/2008/Fahrplan/events/2980.en.html

Rose White

I am a PhD student in sociology at the City University of New York Graduate Center. Starting in spring 2009, I will teach introductory sociology at Baruch College-CUNY. I am also a member of the Graduate Center's doctoral certificate program in Interactive Technology and Pedagogy, where several of us study the history, practice, and theory of interactive media in the classroom.

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-28 | 21:45 CET | 01:00 h | Saal 2 | lecture | Hacking

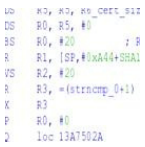

Console Hacking 2008: Wii Fail

Is implementation the enemy of design?

The Nintendo Wii game console has been one of the most popular of all time, selling almost as many units as all of its competitors combined. Despite being cheaper than the PS3 and Xbox360, it contains a sophisticated security architecture that withstood over a year of concerted effort to hack the device. The design itself is impressive; unfortunately, flaws in the implementation (both subtle and severe) render the device easily hacked, with little chance of recovery.

24C3 saw the first public demonstration of unsigned code running on the Wii. A year later, we will present full details of that attack and share the results of another full year of research. We will show the bugs that have been found, the reasons they may have existed, and what attempts the vendor has made to fix them.

Gamers will probably find this talk interesting, but it will be most valuable for anyone who hacks on (or designs) embedded systems. Basic knowledge of crypto is assumed. We will have an area set up in the Hackcenter for those who want to learn more about this subject, before or after the presentation.

http://hackmii.com	*Presenters' blog*
http://wiibrew.org	*Technical info about the Wii*

> http://events.ccc.de/congress/2008/Fahrplan/events/2799.en.html

bushing bushing@hackmii.net

bushing has probably spent a few too many sleepless nights hacking on embedded devices, mostly ARM-based systems. He loves to take things apart, and sometimes to put them back together - but only if they work again afterwards. After a brief demonstration at 24c3, he spent much of 2008 peeling back the layers of security on the Wii. He has degrees in Electrical Engineering and Computer Science.

marcan marcan@marcansoft.com

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-28 | 21:45 CET | 01:00 h | Saal 3 | lecture | Hacking

Tricks: makes you smile

A clever or ingenious device or expedient; adroit technique: the tricks of the trade.

A collection of engaging techniques, some unreleased and some perhaps forgotten, to make pentesting fun again. From layer 3 attacks that still work, to user interaction based exploits that aren't 'clickjacking', to local root privilege escalation without exploits and uncommon web application exploitation techniques.

> http://events.ccc.de/congress/2008/Fahrplan/events/2992.en.html

Francesco `ascii` Ongaro ascii@katamail.com

[redacted] nothing [redacted] to [redacted] hide [redacted]

2008-12-28 | 23:00 CET | 01:00 h | Saal 1 | lecture | Society

Fnord News Show

Wir helfen Euch, die Fnords zu sehen

Das wird dieses Jahr ein apokalyptischer Fnord-Rückblick inklusive Georgien-Krieg und Finanzkrise. Wir versuchen, die Geschehnisse in eine Art Mega-Verschwörungstheorie zu weben, sodass eine Gruppe (oder vielleicht zwei oder drei) an allem schuld sind.

Wir sind uns sicher: Es wird sehr unterhaltsam.

Die Themen werden sein:
• Bankendomino
• Propaganda zum Georgienkrieg
• Terrorismus, Abhören und Schäuble

NICHT der US-Wahlkampf, da hat keiner mehr Bock drauf.

> http://events.ccc.de/congress/2008/Fahrplan/events/2812.en.html

Felix von Leitner

Frank Rieger

2008-12-28 | 23:00 CET | 01:00 h | Saal 2 | lecture | Hacking

Blinkenlights Stereoscope

Behind the scenes of the new light installation

Blinkenlights Stereoscope is the new light installation of Project Blinkenlights, a group that originated form the Chaos Computer Club in 2001. Stereoscope targeted the City Hall in Toronto, Canada and was the biggest and most interactive installation of the group so far. The talk provides insight into how it worked and what technology had been developed to make it all happen.

The Stereoscope project marked a new order of magnitude for the project as the 960 windows and two individual towers of the City Hall in Toronto mean a significant increase in size and complexity compared to the last installation at Bibliothèque nationale de France in Paris. Also the logistics involved in producing the necessary material, shipping it to Canada and have it set up in around two weeks asked for new solutions.

One of the biggest obstacles in quick delivery has been the extensive cabling that is required to control each lamp individually. Although professional light equipment is in widespread use solving a few of these problems, that equipment is very expensive to rent or buy and still does not solve all problems including tons of wires. So the project decided that it's now time for a wireless solution.

Based on the OpenBeacon electronics design, the group developed a new wireless dimmer that can be communicated with over the 2.4 GHz frequency range. Each dimmer controls a single lamp, automatically adjusting to 120V or 240V power source and allowing for displaying up to 16 different shades of light. The whole system was hooked up to a wired backbone with wireless distribution units called Wireless Matrix Control Units (WMCU) that talk to the lamps on one side and listen to a central server on the other.

http://www.blinkenlights.net/stereoscope	*Blinkenlights Stereoscope*
http://www.blinkenlights.net/blog/	*Project Blinkenlights Blog*
http://en.wikipedia.org/wiki/Toronto_City_Hall	*Wikipedia: Toronto City Hall*

> http://events.ccc.de/congress/2008/Fahrplan/events/2934.en.html

Tim Pritlove

Tim Pritlove is a long-time member of the Chaos Computer Club and has been organizing events for the club over a time span of ten year, including the Chaos Communication Congress and the Chaos Communication Camp which he founded. He is also one of the project organizers of Project Blinkenlights. Tim produces the podcast Chaosradio Express and is part of the podcast team for Mobile Macs.

2008-12-28 | 23:00 CET | 01:00 h | Saal 3 | lecture | Science | See paper on p. 245!

Life is a Holodeck!

An overview of holographic techniques

This talk will give you an overview of the different techniques for spacial representation and show you how they work. Starting with a brief history on the invention of stereoscopy and lenticular representation we will quickly get into history and invention of holography, the basic principles and milestones during development through to the latest available applications and technologies. Different types of Holograms will be shown and explained.

• Dennis Gabor - Leith/Upatnieks Invention and the first Hologram 'Train and Bird'
• Holographic Mastering Techniques:
• H1 (classical 2 beam transmission Hologram on a mastering-table)
• Whitelight copies of such H1´s - Lippmann / Denisyuk / Reflection Holograms
• Rainbow/Benton Hologram:
- Photoresist mastering and electroforming of printing plates (shims) for embossed Holograms
- Dot-Matrix Origination - Computer generated Pixel Holograms usually on foil
- True-Color Reflection / Transmission Hologram Mastering
• Video-, Animation and fully computer generated 3D Holograms

Holographic Techniques, Processing, Converting and Special Machinery:
• Microtext + Nanopoints
• 'Hidden' information and 'hidden' Holograms
• Recombining Systems • Electron Beam Mastering
• Electroforming Tanks • Soft- and Hard Embossing Machines: Principle and different types and sizes of conventional holographic embossing machines
• New and future generation manufacturing technology and equipment - Principles of UV/Electron Beam Casting
• Converting Equipment

An overview on market segments as well as current science applications such as HOE´s, (Holographic Optical Elements), holographic interferometry, 3D-projection systems, fully computer generated holograms, optical computing using holographic techniques and other aspects will be given, while at the end we´ll hopefully see what can be done to build your own Holodeck :-)

http://www.holography.ru/techeng.htm Russian Holographer´s lessons

> http://events.ccc.de/congress/2008/Fahrplan/events/3016.en.html

Claus 'HoloClaus' Cohnen cc@hamburg.ccc.de
Claus Cohnen is a heavy-user and developer of holographic technologies since the mid 80ies. He lives in Hamburg, Germany, he is managing director of a company for holographic products,

2008-12-28 | 24:00 CET | 01:00 h | Saal 1 | lecture | Culture

Soviet Unterzoegersdorf

A Nation In Transit

Join a glorious gala presentation with his Excellency and a battalion of members of the Soviet Unterzoegersdorf Military Enforcement Community. We will present the envious 'First World' with the fruits of our techno-labor. Among other triumphs on display will be the second part of an ongoing series of so-called 'Computer Games' or 'Virtual Hyper-Rooms' glorifying the struggles of the Motherland, Soviet Unterzoegersdorf: Sector II. We promise not to mention the SALT II agreement.

'Be as radical as reality.' - Comrade Lenin (1870 - 1924)
People of the world, you, who sit in your superficially secure third countries!

Soviet Unterzoegersdorf is the last existing republic of the USSR. The enclave maintains no diplomatic relationship with the surrounding so-called 'Republic of Austria' or with the Fortress 'European Union'. We persist, undaunted by the downfall of the motherland --- the Soviet Union -- in the early 1990s and its negative effect on our economic situation.

It is a great challenge to secure survival for the small but proud confederation. External reactionary forces put the country in danger. It's a lack of respect due to a morally corrupted and perhaps even non-existing unity of the peoples. The goal of a glorious future is almost unreachable. But his Excellency Ambassador of Soviet Unterzoegersdorf understands the immense transformative power of technology and innovation and how they can improve the lives of all citizens. He realizes that technology offers the tools to prevent and create real change.

So, please join a glorious gala presentation with his Excellency and a battalion of members of the Soviet Unterzoegersdorf Military Enforcement Community. We will present the envious 'First World' with the fruits of our techno-labor. Among other triumphs on display will be the second part of an ongoing series of so-called 'Computer Games' or 'Virtual Hyper-Rooms' glorifying the struggles of the Motherland, Soviet Unterzoegersdorf: Sector II. We promise not to mention the SALT II agreement.

http://www.monochrom.at/suz-game SUZ: The Adventure Game
http://en.wikipedia.org/wiki/Soviet_Unterzoegersdorf Wikipedia entry
http://sovietunterzoegersdorf.metblogs.com/ Soviet Unterzoegersdorf Metroblogging
http://unterzoegersdorf.su Soviet Unterzoegersdorf / unterzoegersdorf.su

> http://events.ccc.de/congress/2008/Fahrplan/events/2773.en.html

grenzfurthner

Johannes Grenzfurthner is an artist, writer, curator, and director. He is the founder of monochrom, an internationally acting art and theory group. He holds a professorship for art theory and art practice at the University of Applied Sciences in Graz, Austria. He is head of the

2008-12-28 | 24:00 CET | 01:00 h | Saal 1 | lecture | Culture

'Arse Elektronika' festival in San Francisco (2007-) and host of 'Roboexotica' (Festival for Cocktail-Robotics) in Vienna (2002-). He also wrote and directed a couple of theater plays. Recurring topics in Johannes' artistic and textual work are contemporary art, activism, performance, humor, philosophy, postmodernism, media theory, cultural studies, popular culture studies, science fiction, and the debate about copyright.

Evelyn Fuerlinger evelyn@monochrom.at
Evelyn Fuerlinger, born 1975 in Austria, lives in Vienna, studied linguistics and literature in Vienna. She is member of the art tech philosophy group monochrom.

Roland Gratzer

Melinda Richka

Michaela Hochrathner

Day 3

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-29 | 11:30 CET | 01:00 h | Saal 1 | lecture | Hacking

Running your own GSM network

This presentation will mark the first public release of a new GPL licensed Free Software project implementing the GSM fixed network, including the various minimal necessary functionality of BSC, MSC, HLR. It will introduce the respective standards and protocols, as well as a short demonstration of an actual phone call between two mobile phones registered to the base station.

On the Ethernet/IP based Internet, we are used to Free Software and general-purpose hardware. The worlds second largest communications network GSM couldn't be any more different. Even though the protocols are standardized and publicly available at the ETSI, all implementations are highly-guarded proprietary secrets of a few major players in the industry. The hardware is even more closed, as there is not a single GSM subscriber or base station chipset with even the least bit of publicly known information.

Nonetheless, in recent years there are a number of different projects working on driving a wedge of Openness into this world. You might have heard about other projects like the THC GSM sniffer project (pure wireshark-like functionality) and OpenBTS (a software defined radio based GSM base station interfacing with the Asterisk VOIP server).

This presentation is about yet another new GSM related Open Source project. A project that follows the GSM specs more closely and actually aims at interoperability with existing equipment such as hardware BTS hooked up via S2M interface to a Linux-running PC.

As part of the presentation we plan to show a live demonstration of a phone call using our own GSM network.

> http://events.ccc.de/congress/2008/Fahrplan/events/3007.en.html

Harald Welte laforge@gnumonks.org

dieter spaar@mirider.augusta.de
Dieter is a self-employed software developer.

2008-12-29 | 11:30 CET | 01:00 h | Saal 2 | lightning | Community

Lightning Talks 2

5 mins of fame

5 minutes for every speaker. Learn about the good, the bad, and the ugly - in software, hardware, projects, and more.

Give a lightning fast talk about your favourite project, program, system - and thereby find people with the same interest to proceed and promote it. Alternatively - give us a good rant about something and give us some good reasons why it should die. ;)

Get right at it, don't waste time by explaining too much, get the main points across, and then let us know how to contact you on the congress for a talk!

Whatever you do - please practise it, and don't be boring. Or else. You have been warned! :-P

> http://events.ccc.de/congress/2008/Fahrplan/events/3047.en.html

SvenG

2008-12-29 | 11:30 CET | 01:00 h | Saal 3 | movie | Culture

Jeder Kontakt hinterlässt Spuren

Die Geschichte des Big Brother Award

Der Film erzählt die Geschichte und Perspektiven der ProtagonistInnen des Big Brother Award. Visuelle Kreisläufe vermitteln einen Eindruck über den Aufwand, die Bedeutung und Wichtigkeit einer Veranstaltung, die jährlich die offensichtlichsten Datenkraken, -händler, -technisierer, -ingenieure, -programmierer und -mißbraucher ... aufs Podest hebt und damit an den Pranger stellt – virtuell wie visuell – konkret wie symbolisch – lustvoll wie ernst.

Filmpremiere: 'Jeder Kontakt hinterlässt Spuren' (HDV, ca. 60 min.)

Der Film vermittelt die Haltung und den Aufwand der einzelnen VertreterInnen der Organsiationen, die den Big Brother Award betreiben. Er erzählt aus der Geschichte und dem sozialen Kontext, in dem der Big Brother Award entstanden ist. Mit Bodenhaftung und Zielsicherheit wird der Prozess der Preisträgerermittlung transparent gemacht und die fabelhafte Stimmung der jährlichen Festveranstaltung mittransportiert. Aus 500 Vorschlägen werden zehn Preisträger ermittelt, und die Bandbreite der Brandmarkung erstreckt sich über eine rassistische CDU-Landesregierung (wie in Hamburg), den Programmierern der nächsten Abrechnungsgeneration für Kfz-Steuern und dem Ausbau bzw. strukturellen Umbau des Überwachungstaates. Alle diese finden in diesem Film ihren Platz – als NEGATIV-Preisträger!

Neben padeluun und Rena Tangens kommen u. a. Rolf Gössner, Karin Schuler (DVD), Alvar C. Freude (Fitug), Frank Rosengart (CCC), Werner Hülsmann (FifF) und der eine oder die andere zu Wort und Tat.

http://www.bigbrotherawards.de *Big Brother Awards*

> http://events.ccc.de/congress/2008/Fahrplan/events/2739.en.html

sven

2008-12-29 | 12:45 CET | 01:00 h | Saal 1 | lecture | Society

eVoting after Nedap and Digital Pen

Why cryptography might not fix the issue of transparent elections

Cryptographic methods have been suggested as a solution of the transparency and auditability issues in electronic voting. This talk introduces some of the suggested approaches and explains why such methods replace one issue with another, rather than fixing it.

Cryptographic methods like Three Ballot, Punchscan, Scantegrity and Bingovoting have been suggested to provide the level of transparency and auditability which is missing in Direct Recording Electronics (DRE), like the NEDAP systems used in Germany's parliamentary elections. These methods introduce a level of complexity into elections which prevents most voters from understanding the election process and its verification. Where elections are currently controlled by the people, trust in the ability of experts is required when cryptographic methods are introduced.

From a more technical perspective, where DRE systems require trust in correct recording and counting of the votes, cryptographic methods might just replace this by the need to trust in the secrecy of the vote.

http://www.punchscan.org/	*Punchscan*
http://people.csail.mit.edu/rivest/Rivest-TheThreeBallotVotingSystem.pdf	*ThreeBallot*
http://www.scantegrity.org/	*Scantegrity*
http://www.bingovoting.de/	*Bingo Voting*

> http://events.ccc.de/congress/2008/Fahrplan/events/3041.en.html

Ulrich Wiesner

2008-12-29 | 12:45 CET | 01:00 h | Saal 2 | lecture | Making

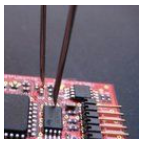

Repurposing the TI EZ430U

with msp430static, solder, and syringe

USB devices are sometimes composed of little more than a microcontroller and a USB device controller. This lecture describes how to reprogram one such device, greatly expanding its potential.

At only twenty dollars, the Texas Instruments EZ430U is a bargain of an in-circuit debugger for the MSP430 microcontroller. The board itself is composed of little more than an MSP430 and a USB to Serial controller. The board's JTAG fuse is unblown, and full schematics are included in public documentation.

This lecture will discuss the use of the EZ430U, not as a debugging tool, but as a development platform in and of itself. Topics will include the writing of replacement firmware, analysis of the default firmware, reprogramming the USB to Serial controller, and potential target applications.

http://travisgoodspeed.blogspot.com/2008/05/repurposing-ti-ez430u-part-1.html
Repurposing the EZ430, Part 1
http://travisgoodspeed.blogspot.com/2008/07/repurposing-ti-ez430u-part-2.html
Repurposing the EZ430, Part 2
http://travisgoodspeed.blogspot.com/2008/08/repurposing-ti-ez430u-part-3.html
Repurposing the EZ430, Part 3

> http://events.ccc.de/congress/2008/Fahrplan/events/2681.en.html

Travis Goodspeed travis@utk.edu
Travis Goodspeed is a neighborly fellow from Knoxville in Southern Appalachia.

2008-12-29 | 12:45 CET | 01:00 h | Saal 3 | lecture | Society

Zehn Big Brother Awards in .at

Rückblick über eine bewegte Zeit

Als erster Big-Brother-Awards-Veranstalter schaffte es Österreich, dieses Jahr die Preise bereits zum zehnten mal zu vergeben. Und obwohl es sich nur um eine Zehnerpotenz handelt, und nicht um eine zur Basis 2, ist es Zeit für einen Rückblick – und einen kleinen Ausblick.

Zehn Veranstaltungen ganz unterschiedlichen Charakters haben wir hinter uns gebracht. Anwaltsbriefe und Klagen erhalten, haarsträubende Geschichten erlebt – und zum Jubiläumsjahr haben wir ein Buch herausgebracht. Aber ein ganz anderes als man vielleicht erwarten würde: ein Science-Fiction-Buch mit 27 Kurzgeschichten – über das, was uns Überwachung bringen könnte. Denn wer beschäftigt sich sonst so intensiv mit Zukunftsszenarien, wie Science Fiction?

Außerdem erzählen wir über Höhen und Tiefen aus zehn Jahren: wo wir einstecken mussten, und wo wir siegreich waren.

http://www.bigbrotherawards.at
http://sf.quintessenz.at

> http://events.ccc.de/congress/2008/Fahrplan/events/2975.en.html

Atr0x

2008-12-29 | 14:00 CET | 01:00 h | Saal 1 | lecture | Society *| See paper on p. 297!*

Neusprech im Überwachungsstaat

Politikersprache zwischen Orwell und Online

Politiker wollen ihre Überwachungspläne schmackhaft machen. Neben der inhaltlichen Verharmlosung von Vorratsdatenspeicherung, Onlinedurchsuchung, Videoüberwachung usw. nutzen sie sprachliche Mittel, um ihre Maßnahmen durchzusetzen. Negativ besetzte Wörter werden durch positive ersetzt und rhetorische Muster werden verwendet, um negative Aspekte auszublenden. Der Vortrag beleuchtet Merkmale der Politikersprache, die in Anlehnung an George Orwell als Neusprech bezeichnet werden kann.

Infolge der Anschläge vom 11. September 2001 ist die „innere Sicherheit“ zu einem wichtigen Thema der Politik geworden. Während sich Politiker durch sicherheitspolitische Maßnahmen Zuspruch erhoffen, ist die mit solchen Maßnahmen verbundene Einschränkung der persönlichen Freiheit problematisch und unpopulär. Daher versuchen Sicherheitspolitiker, ihre Pläne rhetorisch-sprachlich so zu verpacken, dass positive Aspekte hervorgehoben und negative ausgeblendet werden.

http://chaosradio.ccc.de/cre081.html *Chaos Radio Express 081 zum Thema Neusprech*

> http://events.ccc.de/congress/2008/Fahrplan/events/2860.en.html

maha/Martin Haase

Martin Haase ist Professor für Romanische Sprachwissenschaft an der Otto-Friedrich-Universität Bamberg und engagiert sich für freies Wissen und IT-Grundrechte, unter anderem im Chaos Computer Club. Seit 2003 arbeitet er mit bei der Wikipedia und war von 2005 bis 2007 Mitglied im Vorstand von Wikimedia Deutschland e.V., dessen Gründungsmitglied er ist.

2008-12-29 | 14:00 CET | 01:00 h | Saal 2 | lecture | Science | See paper on p. 253!

Privacy in the social semantic web

Social networks based on XMPP

In the last years the static web has moved towards an interactive web – often referred to as the web2.0. People collaboratively write articles in online encyclopedias like Wikipedia or self-portray themselves with profiles in social networks like Myspace. Delicious allows people to tag their bookmarks and share them with friends. Twitter is a short status message service to tell friends what you're doing right now. The diversity of applications attracts a huge amount of users and the application can be used from any computer.

However, many people have privacy concerns with such applications but the advantages and features often outweigh them. Instead of arguing against such services we rather propose an alternative architecture based on the Extensible Messaging and Presence Protocol XMPP.
Within a social network, members can link with each other in order to create a personal network of friends. Often, the number of friends is a kind of "social status" and displayed on a person's profile page. This community aspect attracts a lot of users, especially those who are technically not very experienced. Other social applications don't focus on linking with other members in the first place but allow their users to tag and share special content-types with others. Examples for tagged resources are photos on Flickr, bookmarks on Del.icio.us or publications on Citeulike and Bibsonomy. Both tagging and networking attracted a great deal of attention in the last years.
However, people who want to use the services and share data with others have to provide them to the service maintainer. Most social networks allow to mark data as private or reduce their visibility, but this is not the issue. The main problem we see in current social networks is that private data are given to potentially not trustworthy companies. Users don't know what the companies do with their data or if they can revert their data at all. It may still exist on their servers or in backups. And users can't be sure that private data are always well protected. Security issues often occurred recently in social networks, allowing other to access private data although they were not allowed to. Though the audience of the 25C3 is probably aware of this issues, the technically less experienced people are not. Therefore a simple 'don't use it if you don't like it'-rule is not satisfying. We want to show that technical alternatives to current social networks exist.
We propose a network architecture where users keep the total control of access to their data. Instead of using a client-to-server architecture like traditional social networks do, we use the Extensible Messaging and Presence Protocol XMPP also known as the jabber instance messaging network. Like in instant messenger programs, people can add friends to their personal network. Once they mutually authorized each other, personal data can be exchanged. A public-private-key infrastructure on top of the xmpp communication ensures that message cannot be intercepted or read by any third party – including the xmpp server itself.
The semantic part in our application are the information exchanged between the clients. We decided to use existing ontologies and schemas like FOAF (Friend of a Friend) and the

[redacted] nothing [redacted] to [redacted] hide [redacted]

2008-12-29 | 14:00 CET | 01:00 h | Saal 2 | lecture | Science | See paper on p. 251!

Tag Ontology. In our first prototype users are able to create their personal profile and to bookmark and tag websites. Those data can then be exchanged with friends. Another feature are recursive searches of those bookmarks which allows to retrieve bookmarks of friend-of-friends (as long as they give their permission). We decided to use semantic technologies because we also wanted to show how a semantic web could look like in future. The overall goal is to develop an open, distributed system to exchange information - privately and protected.
The current application is an open source prototype in Java6. The application is available as webstart application and is therefore platform independent. The network is open to other clients and other platform. Other possible applications could be Flash programs, Java applets or browser extensions. Integration into existing instant messenger program is also a possibility.

Ideas for coming features are:
• Integrate the default 'StudiVZ/Facebook' features like pin board, groups, photo-to-person links, etc.,
• Share current location with friends (that is something I would never periodically upload to a website...),
• integrate into local PIM applications: integrated small LDAP server for all address information of friends, RSS feed of latest content from friends,
• OpenID provider (through an HTTP-to-XMPP interface),
• Use the Public-Key in E-Mails, too.

Future Challenges are:
• How can a role-based access control be integrated?
• Once a contact is offline, its information are unavailable. How can they efficiently cached in the network?

My talk will cover some examples of privacy issues and discuss the general architecture. Unless there is concrete interest I won't discuss very research specific topics. I'll give a short introduction into the idea of the semantic web, the arising privacy issues in social networks and the idea of the web-of-trust.

http://www.pace-project.org	*The pace project*
http://www.pace-project.org/jnlp/diki.jnlp	*Java Webstart of the prototype*

> http://events.ccc.de/congress/2008/Fahrplan/events/2873.en.html

Jan Torben
I study Geoinformatics in Münster and write my diploma thesis about – what else could it be – privacy in the social semantic web. This year will be my first visit at the CCC but hopefully not the last.

2008-12-29 | 14:00 CET | 01:00 h | Saal 3 | lecture | Hacking | See paper on p. 149!

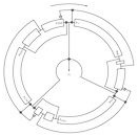

An introduction to new stream cipher designs

Turning data into line noise and back

Even with 'nothing to hide', we want to protect the privacy of our bits and bytes. Encryption is an important tool for this, and stream ciphers are a major class of symmetric-key encryption schemes. Algorithms such as RC4 (used in WEP/WPA, bittorrent, SSL), A5/1 (GSM telephony), E0 (bluetooth), as well as AES in counter (CTR) mode, are important examples of stream ciphers used in everyday applications.

Whereas a block cipher such as AES works by encrypting fixed-length data blocks (and chaining these together in a suitable mode of operation), stream ciphers output an unique and arbitrary-length keystream of pseudorandom bits or bytes, which is simply XORed with the plaintext stream to produce the ciphertext. Advantages of stream ciphers often include smaller hardware footprint and higher encryption speeds than comparable block ciphers such as AES. However, cryptanalysis has led to attacks on many of the existing algorithms.

The ECRYPT Stream Cipher Project (eSTREAM) has been a 4-year project funded by the EU to evaluate new and promising stream ciphers. The project ended in April 2008, with a final portfolio which currently consists of 7 ciphers: 3 suitable for hardware implementation, and 4 aimed at software environments. The portfolio ciphers are considered to provide an advantage over plain AES in at least one significant aspect, but the designs are very different and often suited for different applications.

Since the eSTREAM ciphers are quite new, many of them are not well known outside the academic community. The goal of this talk is to give a very quick presentation of each of the 7 portfolio ciphers: Grain v1, MICKEY v2, Trivium, HC-128, Rabbit, Salsa20/12 and SOSEMANUK.

http://www.ecrypt.eu.org/stream/index.html	*The eSTREAM Project*
http://cr.yp.to/streamciphers.html	*djb's notes on eSTREAM*
http://en.wikipedia.org/wiki/ESTREAM	*Wikipedia has more details*

> http://events.ccc.de/congress/2008/Fahrplan/events/2875.en.html

Tor E. Bjørstad
Tor E. Bjørstad is a PhD student in cryptography at the University of Bergen, Norway.

2008-12-29 | 16:00 CET | 01:00 h | Saal 2 | lecture | Hacking | See paper on p. 209!

The Ultimate Commodore 64 Talk

Everything about the C64 in 64 Minutes

Retrocomputing is cool as never before. People play C64 games in emulators and listen to SID music, but few people know much about the C64 architecture. This talk attempts to communicate 'everything about the C64' to the listener, including its internals and quirks, as well as the tricks that have been used in the demoscene, trying to revive the spirit of times when programmers counted clock cycles and hardware limitations were seen as a challenge.

The Commodore 64 was released in 1982 as an entry- and hobby-level machine competing against the Atari 8 bit series and the Apple II. Compared to other systems on the market, it had a lot of RAM (64 KB), and very sophisticated video and audio hardware. While it was quickly forgotten in the US, it reached its peak in the late 80s in Europe, being a very affordable hobby and game computer. Being the longest running computer of all time, being produced for 12 years, programmers understood the hardware very well, and continued finding new tricks how to create even better graphics effects. 'AGSP' for example, a very sophisticated trick that makes it possible to arbitrarily scroll 'multicolor bitmaps', e.g. for platform games, wasn't used in games until about 1993.
This talk explains all the hardware details of the C64: The programming model of the 6502 CPU family, the Complex Interface Adapters (CIA), the Sound Interface Device, and the programming details as well as common ticks involving the Video Interface Controller (VIC-II). The disk interface will be discussed just as well as the design of the 1541 drive.
The listener will get a good understanding of 8 bit programming and creative programming on extremely limited hardware, as well as common tricks that can be generalized to other systems.

> http://events.ccc.de/congress/2008/Fahrplan/events/2874.en.html

Michael Steil mist@c64.org

Specializing in embedded systems, security systems, operating systems and virtualization, Michael Steil significantly contributed to the Xbox-Linux and GameCube-Linux projects. He holds a Dipl.-Inf. degree from the TU München and is currently employed by a major IT company, working on operating systems kernels.

2008-12-29 | 16:00 CET | 01:00 h | Saal 3 | lecture | Hacking

Hacking into Botnets

Get the real challenge

More and more botnets are coming up every month. The sheer number of hosts infected worldwide is amazing. They collect passwords, send SPAM and DDoS anyone who is annoying the owner. There are two questions coming up immediately:

1) Is there nothing that can be done (besides installing another virus-scanner ;))?
Yes, there are plenty of possibilities. I'd like to present a few in my talk starting from hindering botnets to slowing them down to hijacking bots.

2) Why is nobody doing anything?
The talk is to be seen as an inspiration talk and hopefully leads to lots of discussions.

Did you have SPAM in your mailbox, today? If so, it is very likely that is was sent via a botnet. Maybe such mails are even sent from your computer ... together with your passwords.

There are more and more botnets coming up all the time. Some pass away pretty fast, others grow and grow until they reach sizes of several 10k - 100k hosts worldwide. They are collecting passwords and email-addresses, send tons of SPAM, DDoS anyone who is annoying to the owner, some modify your anti-virus software, others collect personal banking infos and they often can execute any kind of code on the host. Some can be found pretty easily, others include freshest root-kit technologies.

Is there nothing that can be done and if, why is nobody doing anything? (Blackhat Motivation: Wow, there x-thousand hosts running that software that can be remotely controlled. I'd like to get my hands on it (e.g. Massively Distributed Password Cracker). But how?)

The talk is to be an inspiration and start of discussions about techniques (legal and illegal) against botnets. It is divided into two parts. The first and main part is presentation of ideas, research, and real countermeasures against botnets. The second part is about the motivation of people that could but don't do anything.

> http://events.ccc.de/congress/2008/Fahrplan/events/3000.en.html

nano_noname

2008-12-29 | 17:15 CET | 01:00 h | Saal 2 | lecture | Hacking

Security and anonymity vulnerabilities in Tor

Past, present, and future

Now that Germany's data retention law is coming into effect – a few days after the congress – it's time to investigate what it means for the Tor anonymity system, and what changes we need to make, both in terms of system design and in terms of recommendations for relay operators.

Unfortunately, nobody really knows what the data retention law will mean in practice. I'm working with some CCC people to make some good guesses about how it could proceed, and in the talk I'll describe Tor design changes (and culture changes?) that could address each scenario.

http://archives.seul.org/or/dev/Oct-2008/msg00001.html
early outline for my talk material

> http://events.ccc.de/congress/2008/Fahrplan/events/2977.en.html

Roger Dingledine

2008-12-29 | 17:15 CET | 01:00 h | Saal 3 | lecture | Hacking

Squeezing Attack Traces

How to get useable information out of your honeypot

This talk will give an overview about how modern attack analysis tools (dynamic honeypots, an automated shellcode analyzer, and an intrusion signature generator) can be used to get a deep understanding about what attacks do and how they work. A live demo will be given to demonstrate the usage of those tools.

Knowing what's going on in the field of attacks against Internet hosts is one of the most important things for everybody dealing with IT security. People need to stay current with attack technology to understand and implement countermeasures. However, firewall logs and IDS alerts do not provide the details we need. New technologies like honeynets try to bridge this gap: As active sensors they try to catch as much information as possible about an intrusion attempt. But they only collect data most of the time and help little when it comes to actually analyzing attacks.

If we want to understand the attack situation, we need to get some real attack traces first. After that, we can extract the exploit and try to understand, what it does. This can be easy (SQL injection attempts are human readable, for example) but also very hard and time consuming: For a piece of shellcode it would generally be necessary to step over the code in a debugger, a task that is hard to automate. We show a workaround. Finally, once an attack is analyzed, it would be nice to construct a blocking rule or an IDS signature to catch further attempts and prevent other systems from being exploited.

In the talk, we will introduce the idea of using dynamic honeypots for gathering traces of nearly arbitrary server-side attacks. We will show how an automatic shellcode detection and analysis can be performed with a x86 CPU emulation software. Lastly, we will briefly explain how a signature generator can find common parts in different attack traces and how these can be used to assemble a pattern which can be used in a network intrusion detection system.

We will show how to put these tools together in a short live demo.

> http://events.ccc.de/congress/2008/Fahrplan/events/3002.en.html

tw

Georg 'oxff' Wicherski

2008-12-29 | 18:30 CET | 01:00 h | Saal 1 | lecture | Hacking

Attacking NFC mobile phones

First look at the security of NFC mobile phones

Near Field Communication (NFC) based services and mobile phones are starting to appear in the field, therefore it is time to take a look at the security of the services and especially the NFC mobile phones themselves.

The presentation will provide this first look at the security of NFC mobile phones. We will show some known theoretical attacks and how they may work in the field. Further we will present results from analyzing a specific NFC mobile phone, here we will reveal some security issues and methods to exploit them. Also we will provide a small survey of NFC applications in the field. Finally we will release a small set of tools to do further analysis on NFC mobile phones and applications.

If you have an NFC mobile phone I kindly ask you to bring it to the congress and come to me before or after my talk and show it to me, thanks!

http://www.mulliner.org/nfc/

> http://events.ccc.de/congress/2008/Fahrplan/events/2639.en.html

Collin Mulliner collin-25c3@mulliner.org
Security Researcher at day, Hacker at night. Loves gadgets, especially small ones without wires.

2008-12-29 | 18:30 CET | 01:00 h | Saal 2 | lecture | Hacking | See paper on p. 177!

OnionCat – A TOR-based Anonymous VPN

Building an anonymous Internet within the Internet

OnionCat manages to build a complete IP transparent VPN based on TOR's hidden services, provides a simple well-known interface and has the potential to create an anonymous global network which could evolve to a feature- and information-rich network like we know the plain Internet today.

TOR provides so-called 'Hidden Services'. These are services which are location hidden within the TOR network. This means that not only users are hidden but also services (destination). TOR manages this by assigning virtual addresses to them, so-called .onion-URLs. TOR builds all connections based on them.

Unfortunately, access to hidden services is currently not very user-friendly which makes them unattractive although they could provide high privacy in today's world.

OnionCat provides an IP-transparent service which does on-demand connections to designated hidden services. This is a TOR-specific virtual private network (VPN). Because of its IP-transparency any client program can use hidden services without further workarounds.

This talk is about OnionCat in general, gives a brief introduction into its internals and application examples.

http://www.abenteuerland.at/onioncat/ *OnionCat Project Page*
http://www.torproject.org/ *TOR Project home page.*
https://www.torproject.org/doc/design-paper/tor-design.pdf *TOR Designpaper*

> http://events.ccc.de/congress/2008/Fahrplan/events/2828.en.html

rahra
Bernhard has over 20 years system programming experience mainly in development of memory and speed efficient code, parallel computing and networking.

Daniel Haslinger creo@cypherpunk.at
Daniel Haslinger has several years of programming experience in development of flightdata analysis software (
http://fan.rotheneder.com) and network security engineering.

While running his own company dealing with online services, software engineering and network security, he studies 'Information Technology Security' at the University of Applied Science of St.Poelten (Austria). His spare time he spends hunting GPS tagged treasures, electronic and conventional music and personal software projects.

2008-12-29 | 18:30 CET | 01:00 h | Saal 3 | lecture | Hacking | See paper on p. 203!

SWF and the Malware Tragedy

Hide and Seek in A. Flash

This talk rounds up possible web-based attacks using Flash with a particular focus on obfuscation, de-obfuscation and the generic detection of malicious SWF.

While there are some tools out there to analyze AS2 and AS3 based SWF, using various techniques, analysis of SWF can become a nightmare. Starting with a closer look at recent Flash based attacks, this talk will explore ways to recognise these attacks in advance on the one hand, and means to make it even more difficult to prevent them on the other hand. On the way, we will see why and how attackers obfuscate ActionScript code and what methods will probably be used in the future to make detection of malicious payloads much harder.

http://code.google.com/p/erlswf/ *erlswf @ Google Code*
https://www.flashsec.org/ *FlashSec wiki*

> http://events.ccc.de/congress/2008/Fahrplan/events/2596.en.html

BeF

BeF is an enthusiastic open source developer, member of the eventphone phone operation center, ham radio operator and Erlang programmer.

fukami

fukami works for Cologne based web security company called SektionEins and runs a project called FlashSec which is mainly dedicated to Adobe Flash security.

2008-12-29 | 20:30 CET | 01:00 h | Saal 1 | lecture | Hacking

Methods for Understanding Targeted Attacks with Office Documents

As more security features and anti-exploitation mechanisms are added to modern operating systems, attackers are changing their targets to higher-level applications. In the last few years, we have seen increasing targeted attacks using malicious Office documents against both government and non-government entities. These attacks are well publicized in the media; unfortunately, there is not much public information on attack details or exploitation mechanisms employed in the attacks themselves. This presentation aims to fill the gap by offering:

* A brief overview of the Office file format,
* In-depth technical details and practical analytical techniques for triaging and understanding these attacks,
* Defensive mechanisms to reduce the effectiveness of the attacks,
* Forensics evidence that can help trace the attacks,
* Static detection mechanism for these vulnerabilities (i. e., how to write virus signatures for these vulns),
* Information and techniques to help detect these attacks on the wire.

> http://events.ccc.de/congress/2008/Fahrplan/events/2938.en.html

Bruce Dang bda@microsoft.com

2008-12-29 | 20:30 CET | 01:00 h | Saal 2 | lecture | Society

Der elektronische Personalausweis

Endlich wird jeder zum 'Trusted Citizen'

Die Einführung von Fingerabdrücken und biometrischen Gesichtsbildern in den geplanten elektronischen Personalausweis (ePA) ist 2008 beschlossen worden. Versprochen wird uns die sichere Identitätskontrolle, geliefert vom Dienstleister des Vertrauens, der Bundesdruckerei GmbH. Konzeptionelle Fehler aus dem Paßgesetz werden jedoch im neuen Scheckkartenformat des ePA wiederholt.

Der biometrische ePass läßt grüßen: kein Sicherheitszuwachs, aber hohe Kosten auch beim Personalausweis. Biometrie-gestützte Identitätskontrollen werden nun für jedermann verpflichtend, denn anders als noch beim Reisepaß ist mit dem ePA nun jeder Bürger in Deutschland gezwungen, ein biometrisches Ausweisdokument zu beantragen. Damit wird die Lücke im System geschlossen. Da auch dauerhaft hier lebende Ausländer eine elektronische biometrische Karte bekommen, wird also endlich der Traum von der vollerfaßten Bevölkerung wahr.

Die erkennungsdienstliche Behandlung bei ePA ist vorerst für die Fingerabdrücke freiwillig, nicht jedoch für das Gesichtsbild. Was technisch geplant ist und wohin die Reise geht, erzählen wir dem geneigten Zuhörer. Finden Sie sich in der Zwischenzeit schon mal bei ihrer nächstgelegenen biometrischen Registrierstation ein.

> http://events.ccc.de/congress/2008/Fahrplan/events/2895.en.html

46halbe

starbug

2008-12-29 | 21:45 CET | 01:00 h | Saal 1 | lecture | Hacking

Cisco IOS attack and defense

The State of the Art

The talk will cover the past, present and future of Cisco IOS hacking, defense and forensics. Starting from the historic attacks that still work on less well managed parts of the Internet, the powerful common bugs, the classes of binary vulnerabilities and how to exploit them down to the latest methods and techniques, this session will try to give everything in one bag.

To each attack type, we will also see what defensive measures are taken, what should be done and how Cisco forensics people will identify the attack and nail the attacker (or not).

http://www.phenoelit-us.org/ *Phenoelit*
http://cir.recurity-labs.com/ *Cisco Forensics and Wiki*

> http://events.ccc.de/congress/2008/Fahrplan/events/2816.en.html

FX of Phenoelit

2008-12-29 | 21:45 CET | 01:00 h | Saal 2 | lecture | Making

Objects as Software: The Coming Revolution

How RepRap and physical compilers will change the world as we know it (and already have)

How physical compilers (CNC machines, laser cutters, 3D printers, etc) are changing the way we make things, how we think about the nature of objects. This talk will focus on the future of digital manufacturing, and how self-replicating machines will make this technology accessible to everyone: ushering in a new era of technological advance.

Open source software is essentially based on sharing instructions. These coded instructions deal with very tedious things no human would ever want to attempt manually. Luckily, we have computers that understand these instructions and gleefully execute them much faster than any human possibly could. Because of that, we have packets flying across the internet, multi-million pixel displays, and many other amazing technologies. This is all due to computers reliably executing our instructions perfectly and quickly. We know that if we write code on computer A, then it will function exactly the same on computer B (at least in theory ;)

Open source hardware is following a similar route: designs and instructions to create real, physical objects are freely shared. Currently it is fairly difficult to take a design for an object and automatically execute that design as a real object. Due to many exciting developments in the world of open source hardware, this is slowly changing. Things like CNC machines, laser cutters, and 3D printers are becoming more prevalent. These machines will become the physical computers of the coming revolution. As there machines become more and more prevalent, they will also likely increase in quality as well as decrease in cost, allowing people across the globe to digitally share designs which can be created and used locally, without requiring a large skill set to create them. Just as the computer revolution has allowed non-programmers to access the internet and do amazing things with their computers, the physical compiler revolution will allow non-engineers to download and 'print' objects such as robots, appliances, shoes, electronics, and more.

This revolution will transform physical objects into software that be sent around the world in an instant.

http://www.reprap.org/ The RepRap Project

> http://events.ccc.de/congress/2008/Fahrplan/events/2781.en.html

Zach Hoeken hoeken@rrrf.org

Zach Smith is an open source hardware hacker from Brooklyn, NY. He aims to nudge the world in a better direction through awesome technologies.

2008-12-29 | 21:45 CET | 01:00 h | Saal 3 | lecture | Society

Wikileaks

Wikileaks vs. the World

Wikileaks is developing an uncensorable Wikipedia for untraceable mass document leaking and analysis. In the past year, Wikileaks has publicly revealed more sensitive military documents than the entire world's press combined. Its mission has been quite successful after the launch, spawning reportage worldwide and effectively helping to bring about reform on important matters based on factual information. As of now the effort has spawned thousands of press references in major newspapers like The NY Times, The Guardian and the BBC, and tens of thousands in blog posts.

We will talk about experiences that have been made within the first year of its operation, the impact activities on Wikileaks had in various parts of the globe, technical, political and legal challenges faced as well as give an overview of the state of classic and internet media today. We will also talk about conclusions we can derive from these experiences and will present strategies on how investigative journalism, and therefore the fourth estate as the only truly independent control over the state and our future might be resurrected. Lastly we will address why your involvement and that of the technical community is inherently important to ensuring free and uncensored access to information in the future.
During the year of operation we have been able to make many different observations on the state of free information on the internet, the media, governments, military and corporations. We have observed how material that is published is being picked up, sued for, digested, hyped or ignored, and these observations, whether legal, behavioral or qualitative, lead to insights and conclusions that we would like to present and discuss.

Especially we have found the 4th estate as 'the' supposedly independent control over the state, inherently important to any society and its development, is clinically dead, bankrupt and headed in a dangerous direction. While the number of reportage is increasing with bloggers and other new media, the number of genuine reportage, let alone investigative journalism, is rapidly decreasing. This today goes in hand with censorship even in the free world and its media becoming daily routine and increasingly easy. Wikileaks has developed mechanisms that can actively help to address this problem and as has been proven from experience, lead to change and reform. We have found these mechanisms and others in their effectiveness only depend on the awareness and involvement of the public, on our all 'making use of them'.
We want to present these findings in an effort to further this awareness and involvement, especially in the technical community that possesses a lot of the power to shape these important facets of our technologically-driven society, and so in some respect might carry a certain responsibility towards the future of our world.

https://secure.wikileaks.org/wiki/Wikileaks *Wikileaks*

> http://events.ccc.de/congress/2008/Fahrplan/events/2916.en.html

2008-12-29 | 23:00 CET | 02:00 h | Saal 1 | contest | Community

Hacker Jeopardy

Die ultimative Hacker-Quizshow

Das bekannte Quizformat -
aber natürlich mit Themen, die man im Fernsehen nie zu sehen bekäme.

Hacker Jeopardy ist ein Quiz nach dem bekannten umgedrehten Antwort-Frage-Schema. Heise hat es mal 'Zahlenraten für Geeks' genannt, was natürlich eine unfair vereinfachte Darstellung ist – es müssen auch Buchstaben und Sonderzeichen erraten werden. :)

Es werden drei Auswahlrunden gespielt, deren Sieger im Finale gegen den Titelverteidiger des Vorjahres antreten müssen. Wer war das noch?

> http://events.ccc.de/congress/2008/Fahrplan/events/2958.en.html

Stefan 'Sec' Zehl

Ray
Ray's been around for years. Besides lockpicking and random number generating he's interested in linux, electronics and stuff.

2008-12-29 | 23:00 CET | 01:00 h | Saal 2 | lecture | Hacking

We got owned by the (rhymes-with-unease) and didn't even get a lessons learned

life with targeted attacks

The attacks perpetrated by a country in the far east has received numerous mentions, both explicitly and implicitly in the news over the past several years. However relatively few details have emerged as to techniques employed by the intruders. This is largely due to the sensitivity of the selected targets and their will to keep such matters as quiet as possible. As a result of this, large gaps between the various targets ability to cope with the threat have emerged. This talk is the result of being employed as a reverse engineer on an incident response team and illustrates the programs, techniques, network data flow and so on in a level of detail that to the authors knowledge has not been covered else where.

This talk includes a detailed analysis of various malware, such as wauserv.dll, a DLL that turned Windows Update into a beaconing trojan, an analysis of its encraption, the evolution of the techniques employed into this trojan, which evolved into WBC.EXE, and other applications employed. Furthermore, it delves into the network signatures and styles of post-compromise behavior. The presentation is given from my perspective and progresses as the attacks did, from publicly known bugs with public exploits that were unpatched (msjet40.dll) and pwdump.exe to the nearly daily occurence of 0-day attacks that occurred in 2006 in almost all of the various Microsoft Office formats, from embedded executables that were plainly visible in the documents to the ever changing anti-reversing techniques employed.

> http://events.ccc.de/congress/2008/Fahrplan/events/2709.en.html

jf

Day 4

2008-12-30 | 11:30 CET | 01:00 h | Saal 1 | lecture | Community

Why technology sucks

If technology is the solution, politicians are the problem

More and more technology is seen as the ultimate solution for many problems. Lack of understanding and bending rules towards the technology show that politicians and managers have an established level of incompetence. Of course this poses a problem. We tend to forget that hacking also means is having fun with things. Let's ride the incompetence and use technology 'concepts' for the things we want.

We have come to live in a brave new world where technology has added a great value to our existence. The possibilities seem endless and undoubtly are fun. If it isn't fun, breaking it is definitely fun. Living the geeklife is worthwhile.
However managers and politicians have discovered technology too, but they have one serious disability: they aren't geeks and thus by definition not capable of using it in a useful, effective way. Like a spoiled child they won't stop until their toys are installed. In the process of having their way they ignore the real, often damaging effects on society. The aim seems to be to demonstrate how modern and incompetent they really are.
Those 'visionary minds' use technology in a way it wasn't intend and it turns out to be a bad idea. In contrast hackers find fun, unforseen ways of using systems. So when many politicians decry the Netherlands to be an innovative, techsavvy nation, you get a fun talk with a lot of good ideas to respond.

> http://events.ccc.de/congress/2008/Fahrplan/events/3004.en.html

Brenno de Winter brenno@dewinter.com

Brenno De Winter started experimenting with security at the age of 9. He has a background in open source that dates back to 1993 and he contributed to several projects like MySQL, GnuPG, Gnucomo (Gnu Computer Monitoring) and recently started the Small Sister-project for privacy-friendly internet usage. In his daily job he practices security,teaches it and works as an IT-journalist. His writings have triggered several debates in parliament and often raises parlementary questions.

2008-12-30 | 11:30 CET | 01:00 h | Saal 2 | lightning | Community

Lightning Talks 3

5 mins of fame

5 minutes for every speaker. Learn about the good, the bad, and the ugly - in software, hardware, projects, and more.

Give a lightning fast talk about your favourite project, program, system - and thereby find people with the same interest to proceed and promote it. Alternatively - give us a good rant about something and give us some good reasons why it should die. ;)

Get right at it, don't waste time by explaining too much, get the main points across, and then let us know how to contact you on the congress for a talk!

Whatever you do - please practise it, and don't be boring. Or else. You have been warned! :-P

http://www.guckes.net/cccc2006/lightning_talks.html Lightning Talks in 2006

> http://events.ccc.de/congress/2008/Fahrplan/events/2973.en.html

SvenG

2008-12-30 | 11:30 CET | 01:00 h | Saal 3 | lecture | Society

The Privacy Workshop Project

Enhancing the value of privacy in todays students view

The lecture intends to give an overview of the Privacy Workshop project started in Siegen (NRW, Germany) and to animate listeners to participate in the project.

The general idea behind the project "Privacy Protection Workshop Project" is getting in touch with pupils to demonstrate them the importance of privacy protection and demonstrate security technologies that are important in the digital age we live in. The questions we want to raise awareness of, are:

Is there still a need of privacy protection when party videos and pictures in swimwear substitute for job interviews? Who may be able to access my pictures in "SchülerVZ" and other web2.0 platforms? Why is an email more similar to a postcard than a sealed letter? Is it necessary to encode the data on my USB-stick or rather hope it never get lost?
In our experience these important questions are more or less hardly ever considered.

Our workshop-plan distinguishes between two ways of educational methods. First we offer information sessions within the educational curriculum in schools. These will give the pupils an understanding of security measures that are helpful and in our opinion necessary for dealing with SchülerVZ & Co. These sessions are developed in co-operation with their teachers. Second we offer workshops outside the official school hours that will delve more deeply into concepts such as cryptography and secure passwords. The workshop participants are encouraged to put these concepts into direct action by using their brought along USB-sticks and/or Laptops. Due to the complexity of the technologies and their use in day-to-day life we concentrate on Truecrypt and Torpark for USB-sticks and GnuPG for email security.
This workshop will be a new version of: eh2008.koeln.ccc.de/fahrplan/events/2436.de.html
Our first workshop was introduced on EH2008 in March this year. A major part of the presentation was the evaluation of our first workshop focusing on mistakes and how to improve. Promoting our workshop has changed from sending invitations via post to direct and personal contact with teacher and staff at schools. We want to inspire teacher and staff and make them aware of the importance of our project. Currently I am myself a trainee teacher and by now colleagues approach me on a regular basis asking to visit their classes.
Our experience is outlined in an article published December this year in the scholarly journal "Deutschunterricht" published by www.westermann.de the article has been written in collaboration with Axel Krommer: www.deutschdidaktik.ewf.uni-erlangen.de/home/index,id,88,selid,275,type,VAL_MEMO.html

http://www.privacyworkshop.de/ *German Website of the Privacy Workshop Project*

> http://events.ccc.de/congress/2008/Fahrplan/events/2872.en.html

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-30 | 11:30 CET | 01:00 h | Saal 3 | lecture | Society

Christoph Brüning
Student teacher for German and philosophy, FoeBuD activist and CreativeCommons lunatic.

Kai Schubert

2008-12-30 | 12:45 CET | 01:00 h | Saal 2 | lecture | Society | See paper on p. 269!

La Quadrature du Net - Campaigning on Telecoms Package

Pan-european activism for patching a 'pirated' law

La Quadrature du Net (Squaring the Net) is a citizen group informing about legislative projects menacing civil liberties as well as economic and social development in the digital age. Supported by international NGOs (EFF, OSI, ORG, Internautas, Netzwerk Freies Wissen, April, etc.), it aims at providing infrastructure for pan-European activism about such topics as network neutrality, privacy, 'graduated response', etc.

From May to September 24th 2007, a campaign was setup to raise elected representatives', journalists' and public's awareness into the legislative hijack, by the content industries, of the European network regulation law ('Telecoms Package'). A strong mobilization around serious bits of analysis, and proper community tools helped to really influence things.

This talk is about how we got to 'correct' the bad provisions in the text by applying legislative 'patches' in the first reading. If the second reading is happening around the 25C3, then fresh news and updates will be provided.

On the 'graduated response' ('three strikes approach' or 'riposte graduée'), which initiated the founding of La Quadrature du Net, we realized that the industry's strategy for increasing control over the Internet included the use of 'legislative bootstrapping': initiating a law in some country (France in that case), for propagating it on European and worldwide level afterwards.

Thus, using our local expertise on the topic (Olivennes bill about graduated response ordered by N. Sarkozy) could be exported at European level, especially when the 'Telecoms Package' was at the very same moment being hijacked to include IP provisions, to legalize a pan-European graduated response, and to directly harm net neutrality.

La Quadrature du Net was built with the aim of bridging gaps between concerned NGOs across different European countries, providing analysis, pointers, tools and methods allowing everyone to participate on those key issues.

Many good solutions were brought into the text, cleaning the most disturbing parts of it (yet leaving some problematic bits), by constructing dialogues with concerned members of European Parliament (MEPs), producing legal and political analysis, and helping European citizens to participate.

http://www.laquadrature.net La Quadrature du Net
http://www.laquadrature.net/wiki/Telecoms_Package
Wiki page about Telecoms Package
http://www.laquadrat

[redacted] nothing [redacted] to [redacted] hide [redacted]

2008-12-30 | 12:45 CET | 01:00 h | Saal 2 | lecture | Society *| See paper on p. 267!*

http://www.laquadrature.net/wiki/Telecoms_Package_Vote_Sept24_Mobilization
La Quadrature's campaign around first reading of the Package, on Sept. 24th
http://www.laquadrature.net/wiki/Telecoms_package_directives_1st_reading
An analysis and scoring of MEPs recorded votes on Telecoms Package on our 'Political Memory'
http://www.laquadrature.net/wiki/Political_Memory
Political Memory - a tool for tracking members of parliament activity

> http://events.ccc.de/congress/2008/Fahrplan/events/2791.en.html

Jérémie Zimmermann
Co-founder of La Quadrature du Net (Squaring the Net), a citizen group informing about legislative projects menacing civil liberties as well as economic and social development in the digital age.

Markus Beckedahl

2008-12-30 | 12:45 CET | 01:00 h | Saal 3 | lecture | Hacking

Predictable RNG in the vulnerable Debian OpenSSL package

the What and the How

Recently, the Debian project announced an OpenSSL package vulnerability which they had been distributing for the last two years. This bug makes the PRNG predictable, affecting the keys generated by openssl and every other system that uses libssl (eg. openssh, openvpn).

We will talk about this bug (the speaker was the discoverer of this bug), its discovery and publication, its consequences, and exploitation. As well, we will demonstrate some exploitation tools.

> http://events.ccc.de/congress/2008/Fahrplan/events/2995.en.html

Luciano Bello luciano@debian.org
Luciano Bello is an engineer (information systems) and works as a researcher at CITEFA's Si6 Information Security Labs in Buenos Aires, Argentina.

2008-12-30 | 14:00 CET | 01:00 h | Saal 3 | lecture | Culture

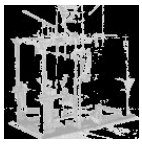

Crafting and Hacking: Separated at Birth

What do hackers have in common with crafters? Lots. While crafting is more often about string and glue than bits and electrons, crafters often feel the same need to create things and manipulate materials into something new. The roots of computing are intertwined with craft around the invention of the Jaquard punchcard loom. We'll look at where the two scenes have gone since then, and what we can gain by reconnecting the hacker world with its softer, more decorative cousin.

In 1801, Joseph Marie Jacquard designed a loom whose patterning was controlled by a piece of perforated pasteboard. This allowed the creation of longer, more complex patterns with a smaller margin for error. Punch cards remained popular through the 70s... but what about craft? Arts and Crafts are rarely considered to be on the forefront of technology... but a number of projects are bringing hackers and crafters back together. Over the last 5 years, the internet has created a huge community of 'open crafting'. Sites like craftster.org encourage people to share ideas and build upon them, much like open source software. It's transformed the 'cottage industry' crafts as well. Crafters go to ecommerce site such as etsy.com or dawanda.com to not only sell their wares, but to participate in communities centered around crafting and building their businesses. Many of the sellers on these sites have no formal business training, it's all DIY.
Bringing these crafting communities online has sparked interest in technology for a lot of crafters. I'll review some really fun projects people have done, like the News Knitter (casualdata.com/newsknitter/), Raph's Twitchy kits (twitchy servo-bots which are super popular with crafters), and some of my own tinkering with the knitting machines (flickr.com/photos/kellbot/2285229844/) and experiments with the laser cutter (www.kellbot.com/2008/09/lasering-rocks/). The current project I'm working on uses Processing to generate patterns for origami boxes, which I then cut on the laser.
My point, and I do have one, is that crafters and hackers have a lot in common. Artists have been setting up shared studios / workspaces, and their needs and interests are similar to those of hackerspaces: space to work, access to equipment, and a supportive community to help them grow their ideas. Our combination craft/hack nights at NYC Resistor have been immensely successful. The two communities have a lot to gain by embracing each other!

> http://events.ccc.de/congress/2008/Fahrplan/events/2777.en.html

Kellbot kellbot@gmail.com
I'm a crafting/hacking nerd living in Brooklyn, NY, USA. My background is in crafting, specifically metalsmithing and jewelry. After moving to New York I got involved with NYC Resistor, who nurtured my geeky tendencies and got me back into programming and electronics. I'm into fuzzy robots, data visualization, and pretty much anything involving NYCR's CO2 laser.

2008-12-30 | 15:15 CET | 01:00 h | Saal 1 | lecture | Hacking

Making the theoretical possible

Attacking a critical piece of Internet infrastructure

> http://events.ccc.de/congress/2008/Fahrplan/events/3023.en.html

Jake

Alexander Sotirov

2008-12-30 | 15:15 CET | 01:00 h | Saal 2 | lecture | Hacking

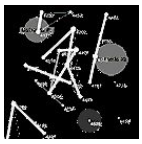

Mining social contacts with active RFID

We describe the implementation of a distributed proximity detection firmware for the OpenBeacon RFID platform. We report on experiments performed during conference gatherings, where the new feature of proximity detection was used to mine and expose patterns of social contact. We discuss some properties of the networks of social contact, and show how these networks can be analyzed, visualized, and used to infer the underlying social structure.

The SocioPatterns project aims to shed light on patterns and statistical regularities in social dynamics. To date, little quantitative information is available about these patterns, and measuring real-world dynamics is indispensable for obtaining a complete picture. In this talk we focus on social contact between people and describe how the OpenBeacon active RFID platform was used to gather experimental data on social contact at a few conference gatherings.

In a variety of contexts, spatial proximity is a good proxy for social interaction. Spatial proximity of persons wearing active RFID tags can be inferred by tracking the location of the tags, and using the position information to decide whether two tags are located nearby. However, locating the tags requires several receiving stations, and contact inference is subjected to errors that limit both its spatial and temporal accuracy. Because of this, we decided to move from contact inference to direct contact detection.

We rewrote the firmware of the OpenBeacon tags specifically targeting proximity detection. We are now able to detect proximity between persons with a very good spatial (~1 m) and temporal (~10 s) resolution. We achieve this by operating the RFID devices in a bi-directional fashion, over multiple radio channels. Tags no longer act as simple beacons, emitting signals for the receiving infrastructure. They exchange messages in a peer-to-peer fashion, to sense their neighborhood and assess contact with other tags. The contact events detected by the RFID network are then relayed to the monitoring infrastructure and post-processed. On suitably tuning the system parameters we achieve reliable detection of face-to-face interaction within about 1 m. This allows, for example, to discriminate who is talking with whom in a small crowded room.

In this talk we discuss our implementation of the contact detection firmware for OpenBeacon tags. We provide some details on data analysis and on the visualization of the longitudinal contact networks we measure. We report the results of an experiment involving about 100 people at a conference, and discuss some interesting statistical regularities of social contact. We also discuss how contact information and trajectory similarity can be used to infer the structure of the social network underlying the community of monitored persons, and how background information can be integrated into this picture. We close by pointing to future directions for research as well as to mashups with social networking services.

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

2008-12-30 | 15:15 CET | 01:00 h | Saal 2 | lecture | Hacking

http://www.sociopatterns.org/	*SocioPatterns*
http://www.openbeacon.org/	*OpenBeacon*
http://www.vimeo.com/1180738	*video 1 on Vimeo*
http://www.vimeo.com/2410580	*video 2 on Vimeo*
http://www.youtube.com/watch?v=ObtVS547lu4	*video 1 on YouTube*
http://www.youtube.com/watch?v=RvFALBMmcz4	*video 2 in YouTube*

> http://events.ccc.de/congress/2008/Fahrplan/events/2899.en.html

Ciro Cattuto

2008-12-30 | 15:15 CET | 01:00 h | Saal 3 | lecture | Making

Pflanzenhacken

Züchten 2.0

Ob Tomaten, Zitronen oder Cannabis: Nutzpflanzen werden längst nicht mehr konventionell in Erde gezüchtet. Von der Auswahl des Saat- und Erbguts bis zur Ernte ist der Anbau von Pflanzen aller Art ein schwieriges, aber spannendes Thema. Die von der Industrie angestellte Forschung hilft auch dem Hobbyzüchter: Pflanzen, die ohne Erde kultiviert und wenige Wochen nach der 'Aussaat' erntereif sind, gehören längst nicht mehr in Science-Fiction-Filme, sondern in den Keller des geneigten Bastlers. Dieser Vortrag soll aufzeigen, dass nicht nur bei Bits'n'Bytes, sondern auch bei Obst und Gemüse durchaus hackbares Potential besteht.

Pflanzenzucht ist keine Esoterik, sondern hackbares Terrain. Zwangsläufig muss der Plantagenbesitzer in spe viel Freude an der Technik mitbringen. So mancher Techie wird sich bei der Konstruktion der eigenen Mikroplantage in seinem Element wiederfinden. Das Thema staubtrocken zu behandeln, wäre witzlos. Während des Vortrags wird es viele Beispiele, Anregungen und Vorführungen zu sehen geben. So gibt dieser Vortrag dem geneigten Gärtner Tipps an die Hand, statt der gekauften Fertiglösung im Ikea-Schrank demnächst aus Heißkleber und Plastikschläuchen entstandene Plantagenträume im eigenen Keller zu beherbergen.

http://eh2008.koeln.ccc.de/fahrplan/events/2448.de.html Vortrag auf dem EH 2008
http://entropia.de/wiki/GPN7:Pflanzen_wachsen_für_Geeks Vortrag auf der GPN7

> http://events.ccc.de/congress/2008/Fahrplan/events/2952.en.html

paul

2008-12-30 | 16:30 CET | 01:30 h | Saal 1 | lecture | Hacking

Security Nightmares 2009

Oder: worüber wir nächstes Jahr lachen werden

Security Nightmares - der jährliche Rückblick auf die IT-Sicherheit und der Security-Glaskugelblick für's nächste Jahr.

Security Nightmares betrachtet die Vergangenheit, Gegenwart und Zukunft von Sicherheitsvorfällen in der IT. Wir machen eine Rückschau auf unsere Vorhersagen der letzten Jahre, unterhalten uns darüber, was sonst noch passiert ist, und wagen dann die Vorschau ins nächste Jahr.

> http://events.ccc.de/congress/2008/Fahrplan/events/3021.en.html

Ron

Frank Rieger

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Chaos Communication Congress
Berlin - bcc, 27.-30.12.2008
http://events.ccc.de/congress/2008/

.....

.....
Papers .
.....

[redacted] nothing [redacted]
[redacted] to [redacted]
[redacted] hide [redacted]

Papers

Hacking

Making

Papers

Hacking

This paper was released February 21, 2008 and published in *Proc. 2008 USENIX Security Symposium*. For the most recent revision, related source code, and videos of demonstration attacks, visit http://citp.princeton.edu/memory.

Lest We Remember: Cold Boot Attacks on Encryption Keys

J. Alex Halderman*, Seth D. Schoen†, Nadia Heninger*, William Clarkson*, William Paul‡, Joseph A. Calandrino*, Ariel J. Feldman*, Jacob Appelbaum, and Edward W. Felten*

*Princeton University †Electronic Frontier Foundation ‡Wind River Systems

{jhalderm, nadiah, wclarkso, jcalandr, ajfeldma, felten}@cs.princeton.edu

schoen@eff.org, wpaul@windriver.com, jacob@appelbaum.net

Abstract

Contrary to popular assumption, DRAMs used in most modern computers retain their contents for several seconds after power is lost, even at room temperature and even if removed from a motherboard. Although DRAMs become less reliable when they are not refreshed, they are not immediately erased, and their contents persist sufficiently for malicious (or forensic) acquisition of usable full-system memory images. We show that this phenomenon limits the ability of an operating system to protect cryptographic key material from an attacker with physical access. We use cold reboots to mount successful attacks on popular disk encryption systems using no special devices or materials. We experimentally characterize the extent and predictability of memory remanence and report that remanence times can be increased dramatically with simple cooling techniques. We offer new algorithms for finding cryptographic keys in memory images and for correcting errors caused by bit decay. Though we discuss several strategies for partially mitigating these risks, we know of no simple remedy that would eliminate them.

1 Introduction

Most security experts assume that a computer's memory is erased almost immediately when it loses power, or that whatever data remains is difficult to retrieve without specialized equipment. We show that these assumptions are incorrect. Ordinary DRAMs typically lose their contents gradually over a period of seconds, even at standard operating temperatures and even if the chips are removed from the motherboard, and data will persist for minutes or even hours if the chips are kept at low temperatures. Residual data can be recovered using simple, nondestructive techniques that require only momentary physical access to the machine.

We present a suite of attacks that exploit DRAM remanence effects to recover cryptographic keys held in memory. They pose a particular threat to laptop users who rely on disk encryption products, since an adversary who steals a laptop while an encrypted disk is mounted could employ our attacks to access the contents, even if the computer is screen-locked or suspended. We demonstrate this risk by defeating several popular disk encryption systems, including BitLocker, TrueCrypt, and FileVault, and we expect many similar products are also vulnerable.

While our principal focus is disk encryption, any sensitive data present in memory when an attacker gains physical access to the system could be subject to attack. Many other security systems are probably vulnerable. For example, we found that Mac OS X leaves the user's login password in memory, where we were able to recover it, and we have constructed attacks for extracting RSA private keys from Apache web servers.

As we discuss in Section 2, certain segments of the computer security and semiconductor physics communities have been conscious of DRAM remanence effects for some time, though strikingly little about them has been published. As a result, many who design, deploy, or rely on secure systems are unaware of these phenomena or the ease with which they can be exploited. To our knowledge, ours is the first comprehensive study of their security consequences.

Highlights and roadmap In Section 3, we describe experiments that we conducted to characterize DRAM remanence in a variety of memory technologies. Contrary to the expectation that DRAM loses its state quickly if it is not regularly refreshed, we found that most DRAM modules retained much of their state without refresh, and even without power, for periods lasting thousands of refresh intervals. At normal operating temperatures, we generally saw a low rate of bit corruption for several seconds, followed by a period of rapid decay. Newer memory technologies, which use higher circuit densities, tended to decay more quickly than older ones. In most cases, we observed that almost all bits decayed at predictable times

and to predictable "ground states" rather than to random values.

We also confirmed that decay rates vary dramatically with temperature. We obtained surface temperatures of approximately $-50\,^{\circ}$C with a simple cooling technique: discharging inverted cans of "canned air" duster spray directly onto the chips. At these temperatures, we typically found that fewer than 1% of bits decayed even after 10 minutes without power. To test the limits of this effect, we submerged DRAM modules in liquid nitrogen (ca. $-196\,^{\circ}$C) and saw decay of only 0.17% after 60 minutes out of the computer.

In Section 4, we present several attacks that exploit DRAM remanence to acquire memory images from which keys and other sensitive data can be extracted. Our attacks come in three variants, of increasing resistance to countermeasures. The simplest is to reboot the machine and launch a custom kernel with a small memory footprint that gives the adversary access to the retained memory. A more advanced attack briefly cuts power to the machine, then restores power and boots a custom kernel; this deprives the operating system of any opportunity to scrub memory before shutting down. An even stronger attack cuts the power and then transplants the DRAM modules to a second PC prepared by the attacker, which extracts their state. This attack additionally deprives the original BIOS and PC hardware of any chance to clear the memory on boot. We have implemented imaging kernels for use with network booting or a USB drive.

If the attacker is forced to cut power to the memory for too long, the data will become corrupted. We propose three methods for reducing corruption and for correcting errors in recovered encryption keys. The first is to cool the memory chips prior to cutting power, which dramatically reduces the error rate. The second is to apply algorithms we have developed for correcting errors in private and symmetric keys. The third is to replicate the physical conditions under which the data was recovered and experimentally measure the decay properties of each memory location; with this information, the attacker can conduct an accelerated error correction procedure. These techniques can be used alone or in combination.

In Section 5, we explore the second error correction method: novel algorithms that can reconstruct cryptographic keys even with relatively high bit-error rates. Rather than attacking the key directly, our methods consider values derived from it, such as key schedules, that provide a higher degree of redundancy. For performance reasons, many applications precompute these values and keep them in memory for as long as the key itself is in use. To reconstruct an AES key, for example, we treat the decayed key schedule as an error correcting code and find the most likely values for the original key. Applying this method to keys with 10% of bits decayed, we can reconstruct nearly any 128-bit AES key within a few seconds. We have devised reconstruction techniques for AES, DES, and RSA keys, and we expect that similar approaches will be possible for other cryptosystems. The vulnerability of precomputation products to such attacks suggests an interesting trade-off between efficiency and security. In Section 6, we present fully automatic techniques for identifying such keys from memory images, even in the presence of bit errors.

We demonstrate the effectiveness of these attacks in Section 7 by attacking several widely used disk encryption products, including BitLocker, TrueCrypt, and FileVault. We have developed a fully automated demonstration attack against BitLocker that allows access to the contents of the disk with only a few minutes of computation. Notably, using BitLocker with a Trusted Platform Module (TPM) sometimes makes it *less* secure, allowing an attacker to gain access to the data even if the machine is stolen while it is completely powered off.

It may be difficult to prevent all the attacks that we describe even with significant changes to the way encryption products are designed and used, but in practice there are a number of safeguards that can provide partial resistance. In Section 8, we suggest a variety of mitigation strategies ranging from methods that average users can apply today to long-term software and hardware changes. Each remedy has limitations and trade-offs. As we conclude in Section 9, it seems there is no simple fix for DRAM remanence vulnerabilities.

Online resources A video demonstration of our attacks and source code for some of our tools are available at http://citp.princeton.edu/memory.

2 Previous Work

Previous researchers have suggested that data in DRAM might survive reboots, and that this fact might have security implications. To our knowledge, however, ours is the first security study to focus on this phenomenon, the first to consider how to reconstruct symmetric keys in the presence of errors, the first to apply such attacks to real disk encryption systems, and the first to offer a systematic discussion of countermeasures.

We owe the suggestion that modern DRAM contents can survive cold boot to Pettersson [33], who seems to have obtained it from Chow, Pfaff, Garfinkel, and Rosenblum [13]. Pettersson suggested that remanence across cold boot could be used to acquire forensic memory images and obtain cryptographic keys, although he did not experiment with the possibility. Chow *et al.* discovered this property in the course of an experiment on data lifetime in running systems. While they did not exploit the

	Memory Type	Chip Maker	Memory Density	Make/Model	Year
A	SDRAM	Infineon	128Mb	Dell Dimension 4100	1999
B	DDR	Samsung	512Mb	Toshiba Portégé	2001
C	DDR	Micron	256Mb	Dell Inspiron 5100	2003
D	DDR2	Infineon	512Mb	IBM T43p	2006
E	DDR2	Elpida	512Mb	IBM x60	2007
F	DDR2	Samsung	512Mb	Lenovo 3000 N100	2007

Table 1: Test systems we used in our experiments

property, they remark on the negative security implications of relying on a reboot to clear memory.

In a recent presentation, MacIver [31] stated that Microsoft considered memory remanence attacks in designing its BitLocker disk encryption system. He acknowledged that BitLocker is vulnerable to having keys extracted by cold-booting a machine when it is used in "basic mode" (where the encrypted disk is mounted automatically without requiring a user to enter any secrets), but he asserted that BitLocker is not vulnerable in "advanced modes" (where a user must provide key material to access the volume). He also discussed cooling memory with dry ice to extend the retention time. MacIver apparently has not published on this subject.

It has been known since the 1970s that DRAM cell contents survive to some extent even at room temperature and that retention times can be increased by cooling. In a 1978 experiment [29], a DRAM showed no data loss for a full week without refresh when cooled with liquid nitrogen. Anderson [2] briefly discusses remanence in his 2001 book:

> [A]n attacker can ... exploit ... memory remanence, the fact that many kinds of computer memory retain some trace of data that have been stored there. ... [M]odern RAM chips exhibit a wide variety of memory remanence behaviors, with the worst of them keeping data for several seconds even at room temperature...

Anderson cites Skorobogatov [40], who found significant data retention times with *static* RAMs at room temperature. Our results for modern DRAMs show even longer retention in some cases.

Anderson's main focus is on "burn-in" effects that occur when data is stored in RAM for an extended period. Gutmann [22, 23] also examines "burn-in," which he attributes to physical changes that occur in semiconductor memories when the same value is stored in a cell for a long time. Accordingly, Gutmann suggests that keys should not be stored in one memory location for longer than several minutes. Our findings concern a different phenomenon: the remanence effects we have studied occur in modern DRAMs even when data is stored only momentarily. These effects do not result from the kind of physical changes that Gutmann described, but rather from the capacitance of DRAM cells.

Other methods for obtaining memory images from live systems include using privileged software running under the host operating system [43], or using DMA transfer on an external bus [19], such as PCI [12], mini-PCI, Firewire [8, 15, 16], or PC Card. Unlike these techniques, our attacks do not require access to a privileged account on the target system, they do not require specialized hardware, and they are resistant to operating system countermeasures.

3 Characterizing Remanence Effects

A DRAM cell is essentially a capacitor. Each cell encodes a single bit by either charging or not charging one of the capacitor's conductors. The other conductor is hard-wired either to power or to ground, depending on the cell's address within the chip [37, 23].

Over time, charge will leak out of the capacitor, and the cell will lose its state or, more precisely, it will decay to its *ground state*, either zero or one depending on whether the fixed conductor of the capacitor is hard-wired to ground or power. To forestall this decay, the cell must be *refreshed*, meaning that the capacitor must be re-charged to hold its value. Specifications for DRAM chips give a *refresh time*, which is the maximum interval that is supposed to pass before a cell is refreshed. The standard refresh time (usually on the order of milliseconds) is meant to achieve extremely high reliability for normal computer operations where even infrequent bit errors could cause serious problems; however, a failure to refresh any individual DRAM cell within this time has only a tiny probability of actually destroying the cell's contents.

We conducted a series of experiments to characterize DRAM remanence effects and better understand the security properties of modern memories. We performed trials using PC systems with different memory technologies, as shown in Table 1. These systems included models from several manufacturers and ranged in age from 9 years to 6 months.

3.1 Decay at operating temperature

Using a modified version of our PXE memory imaging program (see Section 4.1), we filled representative memory regions with a pseudorandom pattern. We read back these memory regions after varying periods of time without refresh and under different temperature conditions, and measured the error rate of each sample. The error rate is the number of bit errors in each sample (the Hamming distance from the pattern we had written) divided by the total number of bits we measured. Since our pseudorandom test pattern contained roughly equal numbers of zeros and ones, we would expect fully decayed memory to have an error rate of approximately 50% .

Our first tests measured the decay rate of each memory module under normal operating temperature, which ranged from 25.5 °C to 44.1 °C, depending on the machine (see Figures 1, 2, and 3). We found that the dimensions of the decay curves varied considerably between machines, with the fastest exhibiting complete data loss in approximately 2.5 seconds and the slowest taking an average of 35 seconds. However, the decay curves all display a similar shape, with an initial period of slow decay, followed by an intermediate period of rapid decay, and then a final period of slow decay.

We calculated best fit curves to the data using the logistic function because MOSFETs, the basic components of a DRAM cell, exhibit a logistic decay curve. We found that machines using newer memory technologies tend to exhibit a shorter time to total decay than machines using older memory technologies, but even the shorter times are long enough to facilitate most of our attacks. We ascribe this trend to the increasing density of the DRAM cells as the technology improves; in general, memory with higher densities have a shorter window where data is recoverable. While this trend might make DRAM retention attacks more difficult in the future, manufacturers also generally seek to *increase* retention times, because DRAMs with long retention require less frequent refresh and have lower power consumption.

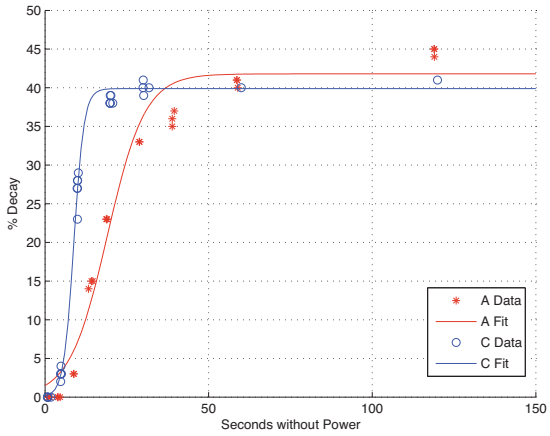

Figure 1: Machines A and C

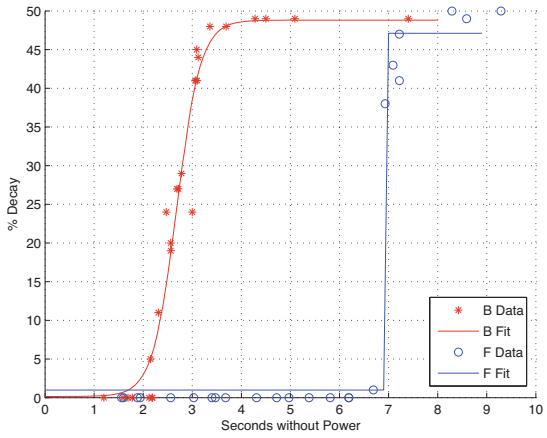

Figure 2: Machines B and F

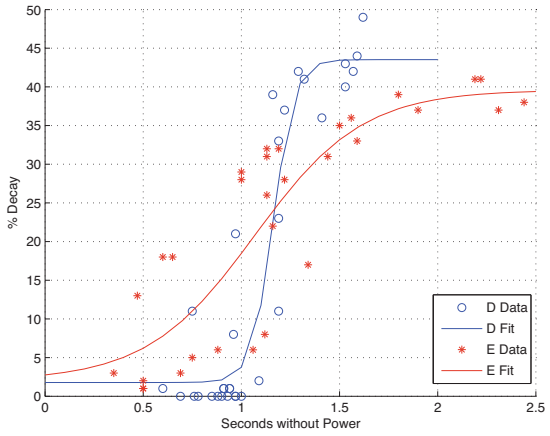

Figure 3: Machines D and E

3.2 Decay at reduced temperature

It has long been known that low temperatures can significantly increase memory devices' retention times [29, 2, 46, 23, 41, 40]. To measure this effect, we performed a second series of tests using machines A–D.

In each trial, we loaded a pseudorandom test pattern into memory, and, with the computer running, cooled the memory module to approximately −50 °C. We then powered off the machine and maintained this temperature until power was restored. We achieved these temperatures using commonly available "canned air" duster products (see Section 4.2), which we discharged, with the can inverted, directly onto the chips.

	Seconds w/o power	Error % at operating temp.	Error % at $-50\,^{\circ}$C
A	60	41	(no errors)
	300	50	0.000095
B	360	50	(no errors)
	600	50	0.000036
C	120	41	0.00105
	360	42	0.00144
D	40	50	0.025
	80	50	0.18

Table 2: Effect of cooling on error rates

As expected, we observed a significantly lower rate of decay under these reduced temperatures (see Table 2). On all of our sample DRAMs, the decay rates were low enough that an attacker who cut power for 60 seconds would recover 99.9% of bits correctly.

As an extreme test of memory cooling, we performed another experiment using liquid nitrogen as an additional cooling agent. We first cooled the memory module of Machine A to $-50\,^{\circ}$C using the "canned air" product. We then cut power to the machine, and quickly removed the DRAM module and placed it in a canister of liquid nitrogen. We kept the memory module submerged in the liquid nitrogen for 60 minutes, then returned it to the machine. We measured only 14,000 bit errors within a 1 MB test region (0.17% decay). This suggests that, even in modern memory modules, data may be recoverable for hours or days with sufficient cooling.

3.3 Decay patterns and predictability

We observed that the DRAMs we studied tended to decay in highly nonuniform patterns. While these patterns varied from chip to chip, they were very predictable in most of the systems we tested. Figure 4 shows the decay in one memory region from Machine A after progressively longer intervals without power.

There seem to be several components to the decay patterns. The most prominent is a gradual decay to the "ground state" as charge leaks out of the memory cells. In the DRAM shown in Figure 4, blocks of cells alternate between a ground state of 0 and a ground state of 1, resulting in the series of horizontal bars. Other DRAM models and other regions within this DRAM exhibited different ground states, depending on how the cells are wired.

We observed a small number of cells that deviated from the "ground state" pattern, possibly due to manufacturing variation. In experiments with 20 or 40 runs, a few "retrograde" cells (typically $\sim 0.05\%$ of memory cells, but larger in a few devices) always decayed to the opposite value of the one predicted by the surrounding ground state pattern. An even smaller number of cells decayed in different directions across runs, with varying probabilities.

Apart from their eventual states, the *order* in which different cells decayed also appeared to be highly predictable. At a fixed temperature, each cell seems to decay after a consistent length of time without power. The relative order in which the cells decayed was largely fixed, even as the decay times were changed by varying the temperature. This may also be a result of manufacturing variations, which result in some cells leaking charge faster than others.

To visualize this effect, we captured degraded memory images, including those shown in Figure 4, after cutting power for intervals ranging from 1 second to 5 minutes, in 1 second increments. We combined the results into a video (available on our web site). Each test interval began with the original image freshly loaded into memory. We might have expected to see a large amount of variation between frames, but instead, most bits appear stable from frame to frame, switching values only once, after the cell's decay interval. The video also shows that the decay intervals themselves follow higher order patterns, likely related to the physical geometry of the DRAM.

3.4 BIOS footprints and memory wiping

Even if memory contents remain intact while power is off, the system BIOS may overwrite portions of memory when the machine boots. In the systems we tested, the BIOS overwrote only relatively small fractions of memory with its own code and data, typically a few megabytes concentrated around the bottom of the address space.

On many machines, the BIOS can perform a destructive memory check during its Power-On Self Test (POST). Most of the machines we examined allowed this test to be disabled or bypassed (sometimes by enabling an option called "Quick Boot").

On other machines, mainly high-end desktops and servers that support ECC memory, we found that the BIOS cleared memory contents without any override option. ECC memory must be set to a known state to avoid spurious errors if memory is read without being initialized [6], and we believe many ECC-capable systems perform this wiping operation whether or not ECC memory is installed.

ECC DRAMs are not immune to retention effects, and an attacker could transfer them to a non-ECC machine that does not wipe its memory on boot. Indeed, ECC memory could turn out to *help* the attacker by making DRAM more resistant to bit errors.

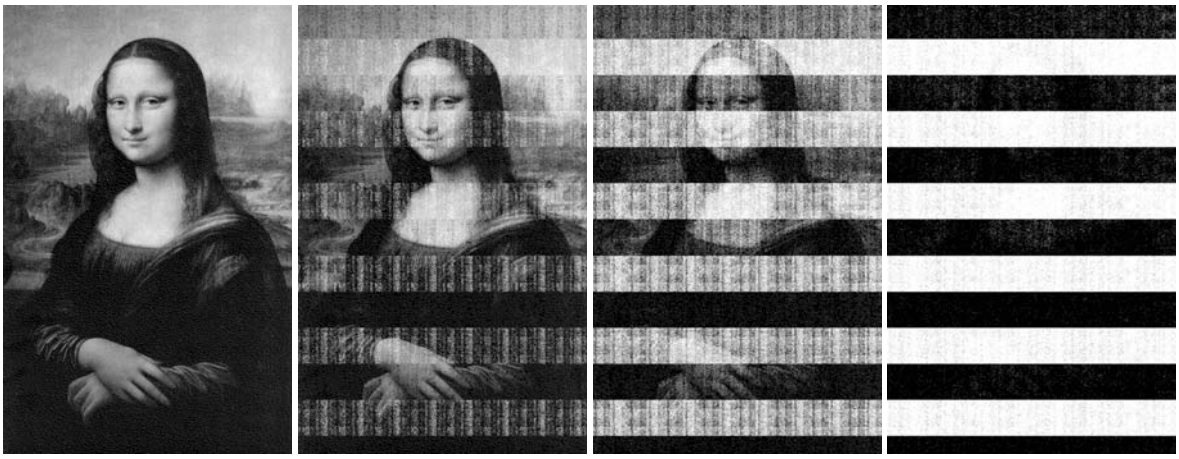

Figure 4: We loaded a bitmap image into memory on Machine A, then cut power for varying lengths of time. After 5 seconds (left), the image is indistinguishable from the original. It gradually becomes more degraded, as shown after 30 seconds, 60 seconds, and 5 minutes.

4 Imaging Residual Memory

Imaging residual memory contents requires no special equipment. When the system boots, the memory controller begins refreshing the DRAM, reading and rewriting each bit value. At this point, the values are fixed, decay halts, and programs running on the system can read any data present using normal memory-access instructions.

4.1 Imaging tools

One challenge is that booting the system will necessarily overwrite some portions of memory. Loading a full operating system would be very destructive. Our approach is to use tiny special-purpose programs that, when booted from either a warm or cold reset state, produce accurate dumps of memory contents to some external medium. These programs use only trivial amounts of RAM, and their memory offsets used can be adjusted to some extent to ensure that data structures of interest are unaffected.

Our memory-imaging tools make use of several different attack vectors to boot a system and extract the contents of its memory. For simplicity, each saves memory images to the medium from which it was booted.

PXE network boot Most modern PCs support network booting via Intel's Preboot Execution Environment (PXE) [25], which provides rudimentary startup and network services. We implemented a tiny (9 KB) standalone application that can be booted via PXE and whose only function is streaming the contents of system RAM via a UDP-based protocol. Since PXE provides a universal API for accessing the underlying network hardware, the same binary image will work unmodified on any PC system with PXE support. In a typical attack setup, a laptop connected to the target machine via an Ethernet crossover cable runs DHCP and TFTP servers as well as a simple client application for receiving the memory data. We have extracted memory images at rates up to 300 Mb/s (around 30 seconds for a 1 GB RAM) with gigabit Ethernet cards.

USB drives Alternatively, most PCs can boot from an external USB device such as a USB hard drive or flash device. We implemented a small (10 KB) plug-in for the SYSLINUX bootloader [3] that can be booted from an external USB device or a regular hard disk. It saves the contents of system RAM into a designated data partition on this device. We succeeded in dumping 1 GB of RAM to a flash drive in approximately 4 minutes.

EFI boot Some recent computers, including all Intel-based Macintosh computers, implement the Extensible Firmware Interface (EFI) instead of a PC BIOS. We have also implemented a memory dumper as an EFI netboot application. We have achieved memory extraction speeds up to 136 Mb/s, and we expect it will be possible to increase this throughput with further optimizations.

iPods We have installed memory imaging tools on an Apple iPod so that it can be used to covertly capture memory dumps without impacting its functionality as a music player. This provides a plausible way to conceal the attack in the wild.

4.2 Imaging attacks

An attacker could use imaging tools like ours in a number of ways, depending on his level of access to the system and the countermeasures employed by hardware and software.

Figure 5: Before powering off the computer, we spray an upside-down canister of multipurpose duster directly onto the memory chips, cooling them to $-50\,^{\circ}$C. At this temperature, the data will persist for several minutes after power loss with minimal error, even if we remove the DIMM from the computer.

Simple reboots The simplest attack is to reboot the machine and configure the BIOS to boot the imaging tool. A warm boot, invoked with the operating system's restart procedure, will normally ensure that the memory has no chance to decay, though software will have an opportunity to wipe sensitive data prior to shutdown. A cold boot, initiated using the system's restart switch or by briefly removing and restoring power, will result in little or no decay depending on the memory's retention time. Restarting the system in this way denies the operating system and applications any chance to scrub memory before shutting down.

Transferring DRAM modules Even if an attacker cannot force a target system to boot memory-imaging tools, or if the target employs countermeasures that erase memory contents during boot, DIMM modules can be physically removed and their contents imaged using another computer selected by the attacker.

Some memory modules exhibit far faster decay than others, but as we discuss in Section 3.2 above, cooling a module before powering it off can slow decay sufficiently to allow it to be transferred to another machine with minimal decay. Widely-available "canned air" dusters, usually containing a compressed fluorohydrocarbon refrigerant, can easily be used for this purpose. When the can is discharged in an inverted position, as shown in Figure 5, it dispenses its contents in liquid form instead of as a gas. The rapid drop in pressure inside the can lowers the temperature of the discharge, and the subsequent evaporation of the refrigerant causes a further chilling. By spraying the contents directly onto memory chips, we can cool their surfaces to $-50\,^{\circ}$C and below. If the DRAM is cooled to this temperature before power is cut and kept cold, we can achieve nearly lossless data recovery even after the chip is out of the computer for several minutes.

Removing the memory modules can also allow the attacker to image memory in address regions where standards BIOSes load their own code during boot. The attacker could remove the primary memory module from the target machine and place it into the secondary DIMM slot (in the same machine or another machine), effectively remapping the data to be imaged into a different part of the address space.

5 Key Reconstruction

Our experiments (see Section 3) show that it is possible to recover memory contents with few bit errors even after cutting power to the system for a brief time, but the presence of even a small amount of error complicates the process of extracting correct cryptographic keys. In this section we present algorithms for correcting errors in symmetric and private keys. These algorithms can correct most errors quickly even in the presence of relatively high bit error probabilities in the range of 5% to 50%, depending on the type of key.

A naïve approach to key error correction is to brute-force search over keys with a low Hamming distance from the decayed key that was retrieved from memory, but this is computationally burdensome even with a moderate amount of unidirectional error. As an example, if only 10% of the ones have decayed to zeros in our memory image, the data recovered from a 256-bit key with an equal number of ones and zeroes has an expected Hamming distance of 12 from the actual key, and the number of such keys is $\binom{128+12}{12} > 2^{56}$.

Our algorithms achieve significantly better performance by considering data other than the raw form of the key. Most encryption programs speed up computation by storing data precomputed from the encryption keys—for block ciphers, this is most often a key schedule, with subkeys for each round; for RSA, this is an extended form of the private key which includes the primes p and q and several other values derived from d. This data contains much more structure than the key itself, and we can use this structure to efficiently reconstruct the original key even in the presence of errors.

These results imply an interesting trade-off between

efficiency and security. All of the disk encryption systems we studied (see Section 7) precompute key schedules and keep them in memory for as long as the encrypted disk is mounted. While this practice saves some computation for each disk block that needs to be encrypted or decrypted, we find that it greatly simplifies key recovery attacks.

Our approach to key reconstruction has the advantage that it is completely self-contained, in that we can recover the key without having to test the decryption of ciphertext. The data derived from the key, and not the decoded plaintext, provides a certificate of the likelihood that we have found the correct key.

We have found it useful to adopt terminology from coding theory. We may imagine that the expanded key schedule forms a sort of *error correcting code* for the key, and the problem of reconstructing a key from memory may be recast as the problem of finding the closest *code word* (valid key schedule) to the data once it has been passed through a channel that has introduced bit errors.

Modeling the decay Our experiments showed that almost all memory bits tend to decay to predictable ground states, with only a tiny fraction flipping in the opposite direction. In describing our algorithms, we assume, for simplicity, that all bits decay to the same ground state. (They can be implemented without this requirement, assuming that the ground state of each bit is known.)

If we assume we have no knowledge of the decay patterns other than the ground state, we can model the decay with the *binary asymmetric channel*, in which the probability of a 1 flipping to 0 is some fixed δ_0 and the probability of a 0 flipping to a 1 is some fixed δ_1.

In practice, the probability of decaying to the ground state approaches 1 as time goes on, while the probability of flipping in the opposite direction remains relatively constant and tiny (less than 0.1% in our tests). The ground state decay probability can be approximated from recovered key data by counting the fraction of 1s and 0s, assuming that the original key data contained roughly equal proportions of each value.

We also observed that bits tended to decay in a predictable order that could be learned over a series of timed decay trials, although the actual order of decay appeared fairly random with respect to location. An attacker with the time and physical access to run such a series of tests could easily adapt any of the approaches in this section to take this order into account and improve the performance of the error-correction. Ideally such tests would be able to replicate the conditions of the memory extraction exactly, but knowledge of the decay order combined with an estimate of the fraction of bit flips is enough to give a very good estimate of an individual decay probability of each bit. This probability can be used in our reconstruction algorithms to prioritize guesses.

For simplicity and generality, we will analyze the algorithms assuming no knowledge of this decay order.

5.1 Reconstructing DES keys

We first apply these methods to develop an error correction technique for DES. The DES key schedule algorithm produces 16 subkeys, each a permutation of a 48-bit subset of bits from the original 56-bit key. Every bit from the original key is repeated in about 14 of the 16 subkeys.

In coding theory terms, we can treat the DES key schedule as a repetition code: the message is a single bit, and the corresponding codeword is a sequence of n copies of this bit. If $\delta_0 = \delta_1 < \frac{1}{2}$, the optimal decoding of such an n-bit codeword is 0 if more than $n/2$ of the recovered bits are 0, and 1 otherwise. For $\delta_0 \neq \delta_1$, the optimal decoding is 0 if more than nr of the recovered bits are 0 and 1 otherwise, where

$$r = \frac{\log(1-\delta_0) - \log\delta_1}{\log(1-\delta_0) + \log(1-\delta_1) - \log\delta_1 - \log\delta_0}.$$

For $\delta_0 = .1$ and $\delta_1 = .001$ (that is, we are in a block with ground state 0), $r = .75$ and this approach will fail to correctly decode a bit only if more than 3 of the 14 copies of a 0 decay to a 1, or more than 11 of the 14 copies of a 1 decay to 0. The probability of this event is less than 10^{-9}. Applying the union bound, the probability that any of the 56 key bits will be incorrectly decoded is at most $56 \times 10^{-9} < 6 \times 10^{-8}$; even at 50% error, the probability that the key can be correctly decoded without resorting to brute force search is more than 98%.

This technique can be trivially extended to correct errors in Triple DES keys. Since Triple DES applies the same key schedule algorithm to two or three 56-bit key components (depending on the version of Triple DES), the probability of correctly decoding each key bit is the same as for regular DES. With a decay rate of $\delta_0 = .5$ and probability $\delta_1 = .001$ of bit flips in the opposite direction, we can correctly decode a 112-bit Triple DES key with at least 97% probability and a 168-bit key with at least 96% probability.

5.2 Reconstructing AES keys

The AES key schedule has a more complex structure than the DES key schedule, but we can still use it to efficiently reconstruct a key in the presence of errors.

A seemingly reasonable approach to this problem would be to search keys in order of distance to the recovered key and output any key whose schedule is sufficiently close to the recovered schedule. Our implementation of this algorithm took twenty minutes to search 10^9 candidate keys in order to reconstruct a key in which 7 zeros

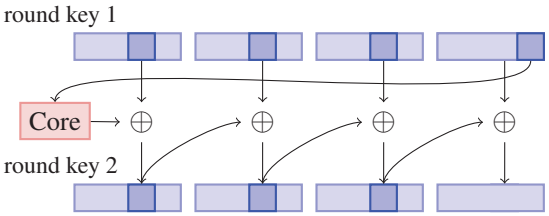

Figure 6: In the 128-bit AES key schedule, three bytes of each round key are entirely determined by four bytes of the preceding round key.

had flipped to ones. At this rate it would take ten days to reconstruct a key with 11 bits flipped.

We can do significantly better by taking advantage of the structure of the AES key schedule. Instead of trying to correct an entire key at once, we can examine a smaller set of bytes at a time. The high amount of linearity in the key schedule is what permits this separability—we can take advantage of pieces that are small enough to brute force optimal decodings for, yet large enough that these decodings are useful to reconstruct the overall key. Once we have a list of possible decodings for these smaller pieces of the key in order of likelihood, we can combine them into a full key to check against the key schedule.

Since each of the decoding steps is quite fast, the running time of the entire algorithm is ultimately limited by the number of combinations we need to check. The number of combinations is still roughly exponential in the number of errors, but it is a vast improvement over brute force searching and is practical in many realistic cases.

Overview of the algorithm For 128-bit keys, an AES key expansion consists of 11 four-word (128-bit) round keys. The first round key is equal to the key itself. Each remaining word of the key schedule is generated either by XORing two words of the key schedule together, or by performing the key schedule core (in which the bytes of a word are rotated and each byte is mapped to a new value) on a word of the key schedule and XORing the result with another word of the key schedule.

Consider a "slice" of the first two round keys consisting of byte i from words 1-3 of the first two round keys, and byte $i-1$ from word 4 of the first round key (as shown in Figure 6). This slice is 7 bytes long, but is uniquely determined by the four bytes from the key. In theory, there are still 2^{32} possibilities to examine for each slice, but we can do quite well by examining them in order of distance to the recovered key. For each possible set of 4 key bytes, we generate the relevant three bytes of the next round key and calculate the probability, given estimates of δ_0 and δ_1, that these seven bytes might have decayed to the corresponding bytes of the recovered round keys.

Now we proceed to guess candidate keys, where a candidate contains a value for each slice of bytes. We consider the candidates in order of decreasing total likelihood as calculated above. For each candidate key we consider, we calculate the expanded key schedule and ask if the likelihood of that expanded key schedule decaying to our recovered key schedule is sufficiently high. If so, then we output the corresponding key as a guess.

When one of δ_0 or δ_1 is very small, this algorithm will almost certainly output a unique guess for the key. To see this, observe that a single bit flipped in the key results in a cascade of bit flips through the key schedule, half of which are likely to flip in the "wrong" direction.

Our implementation of this algorithm is able to reconstruct keys with 15% error (that is, $\delta_0 = .15$ and $\delta_1 = .001$) in a fraction of a second, and about half of keys with 30% error within 30 seconds.

This idea can be extended to 256-bit keys by dividing the words of the key into two sections—words 1–3 and 8, and words 4–7, for example—then comparing the words of the third and fourth round keys generated by the bytes of these words and combining the result into candidate round keys to check.

5.3 Reconstructing tweak keys

The same methods can be applied to reconstruct keys for tweakable encryption modes [30], which are commonly used in disk encryption systems.

LRW LRW augments a block cipher E (and key K_1) by computing a "tweak" X for each data block and encrypting the block using the formula $E_{K_1}(P \oplus X) \oplus X$. A tweak key K_2 is used to compute the tweak, $X = K_2 \otimes I$, where I is the logical block identifier. The operations $\oplus$ and $\otimes$ are performed in the finite field $GF(2^{128})$.

In order to speed tweak computations, implementations commonly precompute multiplication tables of the values $K_2 x^i \mod P$, where x is the primitive element and P is an irreducible polynomial over $GF(2^{128})$ [26]. In practice, $Qx \mod P$ is computed by shifting the bits of Q left by one and possibly XORing with P.

Given a value $K_2 x^i$, we can recover nearly all of the bits of K_2 simply by shifting right by i. The number of bits lost depends on i and the nonzero elements of P. An entire multiplication table will contain many copies of nearly all of the bits of K_2, allowing us to reconstruct the key in much the same way as the DES key schedule.

As an example, we apply this method to reconstruct the LRW key used by the TrueCrypt 4 disk encryption system. TrueCrypt 4 precomputes a 4048-byte multiplication table consisting of 16 blocks of 16 lines of 4 words of 4 bytes each. Line 0 of block 14 contains the tweak key.

The multiplication table is generated line by line from the LRW key by iteratively applying the shift-and-XOR

multiply function to generate four new values, and then XORing all combinations of these four values to create 16 more lines of the table. The shift-and-XOR operation is performed 64 times to generate the table, using the irreducible polynomial $P = x^{128} + x^7 + x^2 + x + 1$. For any of these 64 values, we can shift right i times to recover $128 - (8 + i)$ of the bits of K_2, and use these recovered values to reconstruct K_2 with high probability.

XEX and XTS For XEX [35] and XTS [24] modes, the tweak for block j in sector I is $X = E_{K_2}(I) \otimes x^j$, where I is encrypted with AES and x is the primitive element of $GF(2^{128})$. Assuming the key schedule for K_2 is kept in memory, we can use the AES key reconstruction techniques to reconstruct the tweak key.

5.4 Reconstructing RSA private keys

An RSA public key consists of the modulus N and the public exponent e, while the private key consists of the private exponent d and several optional values: prime factors p and q of N, $d \mod (p-1)$, $d \mod (q-1)$, and $q^{-1} \mod p$. Given N and e, any of the private values is sufficient to generate the others using the Chinese remainder theorem and greatest common divisor algorithms. In practice, RSA implementations store some or all optional values to speed up computation.

There have been a number of results on efficiently reconstructing RSA private keys given a fraction of the bits of private key data. Let $n = \lg N$. N can be factored in polynomial time given the $n/4$ least significant bits of p (Coppersmith [14]), given the $n/4$ least significant bits of d (Boneh, Durfee, and Frankel [9]), or given the $n/4$ least significant bits of $d \mod (p-1)$ (Blömer and May [7]).

These previous results are all based on Coppersmith's method of finding bounded solutions to polynomial equations using lattice basis reduction; the number of contiguous bits recovered from the most or least significant bits of the private key data determines the additive error tolerated in the solution. In our case, the errors may be distributed across all bits of the key data, so we are searching for solutions with low Hamming weight, and these previous approaches do not seem to be directly applicable.

Given the public modulus N and the values $\tilde{p}$ and $\tilde{q}$ recovered from memory, we can deduce values for the original p and q by iteratively reconstructing them from the least-significant bits. For unidirectional decay with probability δ, bits p_i and q_i are uniquely determined by N_i and our guesses for the $i-1$ lower-order bits of p and q (observe that $p_0 = q_0 = 1$), except in the case when $\tilde{p}_i$ and $\tilde{q}_i$ are both in the ground state. This yields a branching process with expected degree $\frac{(3+\delta)^2}{8}$. If decay is not unidirectional, we may use the estimated probabilities to weight the branches at each bit.

Combined with a few heuristics—for example, choose the most likely state first, prune nodes by bounds on the solution, and iteratively increase the bit flips allowed—this results in a practical algorithm for reasonable error rates. This process can likely be improved substantially using additional data recovered from the private key.

We tested an implementation of the algorithm on a fast modern machine. For fifty trials with 1024-bit primes (2048-bit keys) and $\delta = 4\%$, the median reconstruction time was 4.5 seconds. The median number of nodes visited was 16,499, the mean was 621,707, and the standard deviation was 2,136,870. For $\delta = 6\%$, reconstruction required a median of 2.5 minutes, or 227,763 nodes visited.

For 512-bit primes and $\delta = 10\%$, reconstruction required a median of 1 minute, or 188,702 nodes visited.

For larger error rates, we can attempt to reconstruct only the first $n/4$ bits of the key using this process and use the lattice techniques to reconstruct the rest of the key; these computations generally take several hours in practice. For a 1024-bit RSA key, we would need to recover 256 bits of a factor. The expected depth of the tree from our branching reconstruction process would be $(\frac{1}{2} + \delta)^2 256$ (assuming an even distribution of 0s and 1s) and the expected fraction of branches that would need to be examined is $1/2 + \delta^2$.

6 Identifying Keys in Memory

Extracting encryption keys from memory images requires a mechanism for locating the target keys. A simple approach is to test every sequence of bytes to see whether it correctly decrypts some known plaintext. Applying this method to a 1 GB memory image known to contain a 128-bit symmetric key aligned to some 4-byte machine word implies at most 2^{28} possible key values. However, this is only the case if the memory image is perfectly accurate. If there are bit errors in the portion of memory containing the key, the search quickly becomes intractable.

We have developed fully automatic techniques for locating symmetric encryption keys in memory images, even in the presence of bit errors. Our approach is similar to the one we used to correct key bit errors in Section 5. We target the key schedule instead of the key itself, searching for blocks of memory that satisfy (or are close to satisfying) the combinatorial properties of a valid key schedule. Using these methods we have been able to recover keys from closed-source encryption programs without having to disassemble them and reconstruct their key data structures, and we have even recovered partial key schedules that had been overwritten by another program when the memory was reallocated.

Although previous approaches to key recovery do not require a scheduled key to be present in memory, they have other practical drawbacks that limit their usefulness

for our purposes. Shamir and van Someren [39] propose visual and statistical tests of randomness which can quickly identify regions of memory that might contain key material, but their methods are prone to false positives that complicate testing on decayed memory images. Even perfect copies of memory often contain large blocks of random-looking data that might pass these tests (e.g., compressed files). Pettersson [33] suggests a plausibility test for locating a particular program data structure that contains key material based on the range of likely values for each field. This approach requires the operator to manually derive search heuristics for each encryption application, and it is not very robust to memory errors.

6.1 Identifying AES keys

In order to identify scheduled AES keys in a memory image, we propose the following algorithm:

1. Iterate through each byte of memory. Treat the following block of 176 or 240 bytes of memory as an AES key schedule.
2. For each word in the potential key schedule, calculate the Hamming distance from that word to the key schedule word that should have been generated from the surrounding words.
3. If the total number of bits violating the constraints on a correct AES key schedule is sufficiently small, output the key.

We created an application called `keyfind` that implements this algorithm for 128- and 256-bit AES keys. The program takes a memory image as input and outputs a list of likely keys. It assumes that key schedules are contained in contiguous regions of memory and in the byte order used in the AES specification [1]; this can be adjusted to target particular cipher implementations. A threshold parameter controls how many bit errors will be tolerated in candidate key schedules. We apply a quick test of entropy to reduce false positives.

We expect that this approach can be applied to many other ciphers. For example, to identify DES keys based on their key schedule, calculate the distance from each potential subkey to the permutation of the key. A similar method works to identify the precomputed multiplication tables used for advanced cipher modes like LRW (see Section 5.3).

6.2 Identifying RSA keys

Methods proposed for identifying RSA private keys range from the purely algebraic (Shamir and van Someren suggest, for example, multiplying adjacent key-sized blocks of memory [39]) to the *ad hoc* (searching for the RSA Object Identifiers found in ASN.1 key objects [34]). The former ignores the widespread use of standard key formats, and the latter seems insufficiently robust.

The most widely used format for an RSA private key is specified in PKCS #1 [36] as an ASN.1 object of type RSAPrivateKey with the following fields: version, modulus n, publicExponent e, privateExponent d, prime1 p, prime2 q, exponent1 $d \mod (p-1)$, exponent2 $d \mod (q-1)$, coefficient $q^{-1} \mod p$, and optional other information. This object, packaged in DER encoding, is the standard format for storage and interchange of private keys.

This format suggests two techniques we might use for identifying RSA keys in memory: we could search for known *contents* of the fields, or we could look for memory that matches the *structure* of the DER encoding. We tested both of these approaches on a computer running Apache 2.2.3 with `mod_ssl`.

One value in the key object that an attacker is likely to know is the public modulus. (In the case of a web server, the attacker can obtain this and the rest of the public key by querying the server.) We tried searching for the modulus in memory and found several matches, all of them instances of the server's public or private key.

We also tested a key finding method described by Ptacek [34] and others: searching for the RSA Object Identifiers that should mark ASN.1 key objects. This technique yielded only false positives on our test system.

Finally, we experimented with a new method, searching for identifying features of the DER-encoding itself. We looked for the sequence identifier (0x30) followed a few bytes later by the DER encoding of the RSA version number and then by the beginning of the DER encoding of the next field (02 01 00 02). This method found several copies of the server's private key, and no false positives. To locate keys in decayed memory images, we can adapt this technique to search for sequences of bytes with low Hamming distance to these markers and check that the subsequent bytes satisfy some heuristic entropy bound.

7 Attacking Encrypted Disks

Encrypting hard drives is an increasingly common countermeasure against data theft, and many users assume that disk encryption products will protect sensitive data even if an attacker has physical access to the machine. A California law adopted in 2002 [10] requires disclosure of possible compromises of personal information, but offers a safe harbor whenever data was "encrypted." Though the law does not include any specific technical standards, many observers have recommended the use of full-disk or file system encryption to obtain the benefit of this safe harbor. (At least 38 other states have enacted data breach notification legislation [32].) Our results below suggest

that disk encryption, while valuable, is not necessarily a sufficient defense. We find that a moderately skilled attacker can circumvent many widely used disk encryption products if a laptop is stolen while it is powered on or suspended.

We have applied some of the tools developed in this paper to attack popular on-the-fly disk encryption systems. The most time-consuming parts of these tests were generally developing system-specific attacks and setting up the encrypted disks. Actually imaging memory and locating keys took only a few minutes and were almost fully automated by our tools. We expect that most disk encryption systems are vulnerable to such attacks.

BitLocker BitLocker, which is included with some versions of Windows Vista, operates as a filter driver that resides between the file system and the disk driver, encrypting and decrypting individual sectors on demand. The keys used to encrypt the disk reside in RAM, in scheduled form, for as long as the disk is mounted.

In a paper released by Microsoft, Ferguson [21] describes BitLocker in enough detail both to discover the roles of the various keys and to program an independent implementation of the BitLocker encryption algorithm without reverse engineering any software. BitLocker uses the same pair of AES keys to encrypt every sector on the disk: a sector pad key and a CBC encryption key. These keys are, in turn, indirectly encrypted by the disk's master key. To encrypt a sector, the plaintext is first XORed with a pad generated by encrypting the byte offset of the sector under the sector pad key. Next, the data is fed through two diffuser functions, which use a Microsoft-developed algorithm called Elephant. The purpose of these un-keyed functions is solely to increase the probability that modifications to any bits of the ciphertext will cause unpredictable modifications to the entire plaintext sector. Finally, the data is encrypted using AES in CBC mode using the CBC encryption key. The initialization vector is computed by encrypting the byte offset of the sector under the CBC encryption key.

We have created a fully-automated demonstration attack called BitUnlocker. It consists of an external USB hard disk containing Linux, a custom SYSLINUX-based bootloader, and a FUSD [20] filter driver that allows BitLocker volumes to be mounted under Linux. To use BitUnlocker, one first cuts the power to a running Windows Vista system, connects the USB disk, and then reboots the system off of the external drive. BitUnlocker then automatically dumps the memory image to the external disk, runs `keyfind` on the image to determine candidate keys, tries all combinations of the candidates (for the sector pad key and the CBC encryption key), and, if the correct keys are found, mounts the BitLocker encrypted volume. Once the encrypted volume has been mounted, one can browse it like any other volume in Linux. On a modern laptop with 2 GB of RAM, we found that this entire process took approximately 25 minutes.

BitLocker differs from other disk encryption products in the way that it protects the keys when the disk is not mounted. In its default "basic mode," BitLocker protects the disk's master key solely with the Trusted Platform Module (TPM) found on many modern PCs. This configuration, which may be quite widely used [21], is particularly vulnerable to our attack, because the disk encryption keys can be extracted with our attacks even if the computer is powered off for a long time. When the machine boots, the keys will be loaded into RAM automatically (before the login screen) without the entry of any secrets.

It appears that Microsoft is aware of this problem [31] and recommends configuring BitLocker in "advanced mode," where it protects the disk key using the TPM along with a password or a key on a removable USB device. However, even with these measures, BitLocker is vulnerable if an attacker gets to the system while the screen is locked or the computer is asleep (though not if it is hibernating or powered off).

FileVault Apple's FileVault disk encryption software has been examined and reverse-engineered in some detail [44]. In Mac OS X 10.4, FileVault uses 128-bit AES in CBC mode. A user-supplied password decrypts a header that contains both the AES key and a second key k_2 used to compute IVs. The IV for a disk block with logical index I is computed as HMAC-SHA1$_{k_2}(I)$.

We used our EFI memory imaging program to extract a memory image from an Intel-based Macintosh system with a FileVault volume mounted. Our `keyfind` program automatically identified the FileVault AES key, which did not contain any bit errors in our tests.

With the recovered AES key but not the IV key, we can decrypt 4080 bytes of each 4096 byte disk block (all except the first AES block). The IV key is present in memory. Assuming no bits in the IV key decay, an attacker can identify it by testing all 160-bit substrings of memory to see whether they create a plausible plaintext when XORed with the decryption of the first part of the disk block. The AES and IV keys together allow full decryption of the volume using programs like `vilefault` [45].

In the process of testing FileVault, we discovered that Mac OS X 10.4 and 10.5 keep multiple copies of the user's login password in memory, where they are vulnerable to imaging attacks. Login passwords are often used to protect the default keychain, which may protect passphrases for FileVault disk images.

TrueCrypt TrueCrypt is a popular open-source disk encryption product for the Windows, Mac OS, and Linux platforms. It supports a variety of ciphers, including AES, Serpent, and Twofish. In version 4, all ciphers used LRW mode; in version 5, they use XTS mode (see Section 5.3).

TrueCrypt stores a cipher key and a tweak key in the volume header for each disk, which is then encrypted with a separate key derived from a user-entered password.

We tested TrueCrypt versions 4.3a and 5.0a running on a Linux system. We mounted a volume encrypted with a 256-bit AES key, then briefly cut power to the system and used our memory imaging tools to record an image of the retained memory data. In both cases, our `keyfind` program was able to identify the 256-bit AES encryption key, which did not contain any bit errors. For TrueCrypt 5.0a, `keyfind` was also able to recover the 256-bit AES XTS tweak key without errors.

To decrypt TrueCrypt 4 disks, we also need the LRW tweak key. We observed that TrueCrypt 4 stores the LRW key in the four words immediately preceding the AES key schedule. In our test memory image, the LRW key did not contain any bit errors. (Had errors occurred, we could have recovered the correct key by applying the techniques we developed in Section 5.3.)

dm-crypt Linux kernels starting with 2.6 include built-in support for dm-crypt, an on-the-fly disk encryption subsystem. The dm-crypt subsystem handles a variety of ciphers and modes, but defaults to 128-bit AES in CBC mode with non-keyed IVs.

We tested a dm-crypt volume created and mounted using the LUKS (Linux Unified Key Setup) branch of the `cryptsetup` utility and kernel version 2.6.20. The volume used the default AES-CBC format. We briefly powered down the system and captured a memory image with our PXE kernel. Our `keyfind` program identified the correct 128-bit AES key, which did not contain any bit errors. After recovering this key, an attacker could decrypt and mount the dm-crypt volume by modifying the `cryptsetup` program to allow input of the raw key.

Loop-AES Loop-AES is an on-the-fly disk encryption package for Linux systems. In its recommended configuration, it uses a so-called "multi-key-v3" encryption mode, in which each disk block is encrypted with one of 64 encryption keys. By default, it encrypts sectors with AES in CBC mode, using an additional AES key to generate IVs.

We configured an encrypted disk with Loop-AES version 3.2b using 128-bit AES encryption in "multi-key-v3" mode. After imaging the contents of RAM, we applied our `keyfind` program, which revealed the 65 AES keys. An attacker could identify which of these keys correspond to which encrypted disk blocks by performing a series of trial decryptions. Then, the attacker could modify the Linux `losetup` utility to mount the encrypted disk with the recovered keys.

Loop-AES attempts to guard against the long-term memory burn-in effects described by Gutmann [23] and others. For each of the 65 AES keys, it maintains two copies of the key schedule in memory, one normal copy and one with each bit inverted. It periodically swaps these copies, ensuring that every memory cell stores a 0 bit for as much time as it stores a 1 bit. Not only does this fail to prevent the memory remanence attacks that we describe here, but it also makes it easier to identify which keys belong to Loop-AES and to recover the keys in the presence of memory errors. After recovering the regular AES key schedules using a program like `keyfind`, the attacker can search the memory image for the inverted key schedules. Since very few programs maintain both regular and inverted key schedules in this way, those keys are highly likely to belong to Loop-AES. Having two related copies of each key schedule provides additional redundancy that can be used to identify which bit positions are likely to contain errors.

8 Countermeasures and their Limitations

Memory imaging attacks are difficult to defend against because cryptographic keys that are in active use need to be stored *somewhere*. Our suggested countermeasures focus on discarding or obscuring encryption keys before an adversary might gain physical access, preventing memory-dumping software from being executed on the machine, physically protecting DRAM chips, and possibly making the contents of memory decay more readily.

Scrubbing memory Countermeasures begin with efforts to avoid storing keys in memory. Software should overwrite keys when they are no longer needed, and it should attempt to prevent keys from being paged to disk. Runtime libraries and operating systems should clear memory proactively; Chow *et al.* show that this precaution need not be expensive [13]. Of course, these precautions cannot protect keys that must be kept in memory because they are still in use, such as the keys used by encrypted disks or secure web servers.

Systems can also clear memory at boot time. Some PCs can be configured to clear RAM at startup via a destructive Power-On Self-Test (POST) before they attempt to load an operating system. If the attacker cannot bypass the POST, he cannot image the PC's memory with locally-executing software, though he could still physically move the memory chips to different computer with a more permissive BIOS.

Limiting booting from network or removable media Many of our attacks involve booting a system via the network or from removable media. Computers can be configured to require an administrative password to boot from these sources. We note, however, that even if a system will boot only from the primary hard drive, an attacker could still swap out this drive, or, in many cases,

reset the computer's NVRAM to re-enable booting from removable media.

Suspending a system safely Our results show that simply locking the screen of a computer (i.e., keeping the system running but requiring entry of a password before the system will interact with the user) does not protect the contents of memory. Suspending a laptop's state ("sleeping") is also ineffective, even if the machine enters screen-lock on awakening, since an adversary could simply awaken the laptop, power-cycle it, and then extract its memory state. Suspending-to-disk ("hibernating") may also be ineffective unless an externally-held secret is required to resume normal operations.

With most disk encryption systems, users can protect themselves by powering off the machine completely when it is not in use. (BitLocker in "basic" TPM mode remains vulnerable, since the system will automatically mount the disk when the machine is powered on.) Memory contents may be retained for a short period, so the owner should guard the machine for a minute or so after removing power. Though effective, this countermeasure is inconvenient, since the user will have to wait through the lengthy boot process before accessing the machine again.

Suspending can be made safe by requiring a password or other external secret to reawaken the machine, and encrypting the contents of memory under a key derived from the password. The password must be strong (or strengthened), as an attacker who can extract memory contents can then try an offline password-guessing attack. If encrypting all of memory is too expensive, the system could encrypt only those pages or regions containing important keys. Some existing systems can be configured to suspend safely in this sense, although this is often not the default behavior [5].

Avoiding precomputation Our attacks show that using precomputation to speed cryptographic operations can make keys more vulnerable. Precomputation tends to lead to redundant storage of key information, which can help an attacker reconstruct keys in the presence of bit errors, as described in Section 5.

Avoiding precomputation may hurt performance, as potentially expensive computations will be repeated. (Disk encryption systems are often implemented on top of OS- and drive-level caches, so they are more performance-sensitive than might be assumed.) Compromises are possible; for example, precomputed values could be cached for a predetermined period of time and discarded if not re-used within that interval. This approach accepts some vulnerability in exchange for reducing computation, a sensible tradeoff in some situations.

Key expansion Another defense against key reconstruction is to apply some transform to the key as it is stored in memory in order to make it more difficult to reconstruct in the case of errors. This problem has been considered from a theoretical perspective; Canetti *et al.* [11] define the notion of an *exposure-resilient function* whose input remains secret even if all but some small fraction of the output is revealed, and show that the existence of this primitive is equivalent to the existence of one-way functions.

In practice, suppose we have a key K which is not currently in use but will be needed later. We cannot overwrite the key but we want to make it more resistant to reconstruction. One way to do this is to allocate a large B-bit buffer, fill the buffer with random data R, then store $K \oplus H(R)$ where H is a hash function such as SHA-256.

Now suppose there is a power-cutting attack which causes d of the bits in this buffer to be flipped. If the hash function is strong, the adversary must search a space of size $\binom{B/2+d}{d}$ to discover which bits were flipped of the roughly $B/2$ that could have decayed. If B is large, this search will be prohibitive even when d is relatively small.

In principle, all keys could be stored in this way, recomputed when in use, and deleted immediately after. Alternatively, we could sometimes keep keys in memory, introducing the precomputation tradeoff discussed above.

For greater protection, the operating system could perform tests to identify memory locations that are especially quick to decay, and use these to store key material.

Physical defenses Some of our attacks rely on physical access to DRAM chips or modules. These attacks can be mitigated by physically protecting the memory. For example, DRAM modules could be locked in place inside the machine, or encased in a material such as epoxy to frustrate attempts to remove or access them. Similarly, the system could respond to low temperatures or opening of the computer's case by attempting to overwrite memory, although these defenses require active sensor systems with their own backup power supply. Many of these techniques are associated with specialized tamper-resistant hardware such as the IBM 4758 coprocessor [18, 41] and could add considerable cost to a PC. However, a small amount of memory soldered to a motherboard could be added at relatively low cost.

Architectural changes Some countermeasures try to change the machine's architecture. This will not help on existing machines, but it might make future machines more secure.

One approach is to find or design DRAM systems that lose their state quickly. This might be difficult, given the tension between the desire to make memory decay quickly and the desire to keep the probability of decay within a DRAM refresh interval vanishingly small.

Another approach is to add key-store hardware that erases its state on power-up, reset, and shutdown. This would provide a safe place to put a few keys, though precomputation of derived keys would still pose a risk.

Others have proposed architectures that would routinely encrypt the contents of memory for security purposes [28, 27, 17]. These would apparently prevent the attacks we describe, as long as the encryption keys were destroyed on reset or power loss.

Encrypting in the disk controller Another approach is to encrypt data in the hard disk controller hardware, as in Full Disk Encryption (FDE) systems such as Seagate's "DriveTrust" technology [38].

In its basic form, this approach uses a write-only *key register* in the disk controller, into which the software can write a symmetric encryption key. Data blocks are encrypted, using the key from the key register, before writing to the disk. Similarly, blocks are decrypted after reading. This allows encrypted storage of all blocks on a disk, without any software modifications beyond what is required to initialize the key register.

This approach differs from typical disk encryption systems in that encryption and decryption are done by the disk controller rather than by software in the main CPU, and that the main encryption keys are stored in the disk controller rather than in DRAM.

To be secure, such a system must ensure that the key register is erased whenever a new operating system is booted on the computer; otherwise, an attacker can reboot into a malicious kernel that simply reads the disk contents. For similar reasons, the system must also ensure that the key register is erased whenever an attacker attempts to move the disk controller to another computer (even if the attacker maintains power while doing so).

Some systems built more sophisticated APIs, implemented by software on the disk controller, on top of this basic facility. Such APIs, and their implementation, would require further security analyses.

We have not evaluated any specific systems of this type. We leave such analyses for future work.

Trusted computing Trusted Computing hardware, in the form of Trusted Platform Modules (TPMs) [42] is now deployed in some personal computers. Though useful against some attacks, today's Trusted Computing hardware does not seem to prevent the attacks described here.

Deployed TCG TPMs do not implement bulk encryption. Instead, they monitor boot history in order to decide (or help other machines decide) whether it is safe to store a key in RAM. If a software module wants to use a key, it can arrange that the usable form of that key will not be stored in RAM unless the boot process has gone as expected [31]. However, once the key is stored in RAM, it is subject to our attacks. TPMs can prevent a key from being loaded into memory for use, but they cannot prevent it from being captured once it is in memory.

9 Conclusions

Contrary to popular belief, DRAMs hold their values for surprisingly long intervals without power or refresh. Our experiments show that this fact enables a variety of security attacks that can extract sensitive information such as cryptographic keys from memory, despite the operating system's efforts to protect memory contents. The attacks we describe are practical—for example, we have used them to defeat several popular disk encryption systems.

Other types of software may be similarly vulnerable. DRM systems often rely on symmetric keys stored in memory, which may be recoverable using the techniques outlined in our paper. As we have shown, SSL-enabled web servers are vulnerable, since they often keep in memory private keys needed to establish SSL sessions. Furthermore, methods similar to our key-finder would likely be effective for locating passwords, account numbers, or other sensitive data in memory.

There seems to be no easy remedy for these vulnerabilities. Simple software changes have benefits and drawbacks; hardware changes are possible but will require time and expense; and today's Trusted Computing technologies cannot protect keys that are already in memory. The risk seems highest for laptops, which are often taken out in public in states that are vulnerable to our attacks. These risks imply that disk encryption on laptops, while beneficial, does not guarantee protection.

Ultimately, it might become necessary to treat DRAM as untrusted, and to avoid storing sensitive data there, but this will not be feasible until architectures are changed to give software a safe place to keep its keys.

Acknowledgments

We thank Andrew Appel, Jesse Burns, Grey David, Laura Felten, Christian Fromme, Dan Good, Peter Gutmann, Benjamin Mako Hill, David Hulton, Brie Ilenda, Scott Karlin, David Molnar, Tim Newsham, Chris Palmer, Audrey Penven, David Robinson, Kragen Sitaker, N.J.A. Sloane, Gregory Sutter, Sam Taylor, Ralf-Philipp Weinmann, and Bill Zeller for their helpful comments, suggestions, and contributions. Without them, this paper would not have been possible.

This material is based in part upon work supported under a National Science Foundation Graduate Research Fellowship. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Calandrino performed this research under an appointment to the Department of Homeland Security (DHS) Scholarship and Fellowship Program, administered by the Oak Ridge Institute for Science and Education (ORISE) through an interagency agreement between the U.S. Department of Energy (DOE) and DHS. ORISE is managed by Oak Ridge Associated Universities (ORAU) under DOE contract number DE-AC05-06OR23100.

All opinions expressed in this paper are the author's and do not necessarily reflect the policies and views of DHS, DOE, or ORAU/ORISE.

References

[1] Advanced Encryption Standard. National Institute of Standards and Technology, FIPS-197, Nov. 2001.

[2] ANDERSON, R. *Security Engineering: A Guide to Building Dependable Distributed Systems*, first ed. Wiley, Jan. 2001, p. 282.

[3] ANVIN, H. P. SYSLINUX. http://syslinux.zytor.com/.

[4] ARBAUGH, W., FARBER, D., AND SMITH, J. A secure and reliable bootstrap architecture. In *Proc. IEEE Symp. on Security and Privacy* (May 1997), pp. 65–71.

[5] BAR-LEV, A. Linux, Loop-AES and optional smartcard based disk encryption. http://wiki.tuxonice.net/EncryptedSwapAndRoot, Nov. 2007.

[6] BARRY, P., AND HARTNETT, G. *Designing Embedded Networking Applications: Essential Insights for Developers of Intel IXP4XX Network Processor Systems*, first ed. Intel Press, May 2005, p. 47.

[7] BLÖMER, J., AND MAY, A. New partial key exposure attacks on RSA. In *Proc. CRYPTO 2003* (2003), pp. 27–43.

[8] BOILEAU, A. Hit by a bus: Physical access attacks with Firewire. Presentation, Ruxcon, 2006.

[9] BONEH, D., DURFEE, G., AND FRANKEL, Y. An attack on RSA given a small fraction of the private key bits. In *Advances in Cryptology – ASIACRYPT '98* (1998), pp. 25–34.

[10] CALIFORNIA STATUTES. Cal. Civ. Code §1798.82, created by S.B. 1386, Aug. 2002.

[11] CANETTI, R., DODIS, Y., HALEVI, S., KUSHILEVITZ, E., AND SAHAI, A. Exposure-resilient functions and all-or-nothing transforms. In *Advances in Cryptology – EUROCRYPT 2000* (2000), vol. 1807/2000, pp. 453–469.

[12] CARRIER, B. D., AND GRAND, J. A hardware-based memory acquisition procedure for digital investigations. *Digital Investigation 1* (Dec. 2003), 50–60.

[13] CHOW, J., PFAFF, B., GARFINKEL, T., AND ROSENBLUM, M. Shredding your garbage: Reducing data lifetime through secure deallocation. In *Proc. 14th USENIX Security Symposium* (Aug. 2005), pp. 331–346.

[14] COPPERSMITH, D. Small solutions to polynomial equations, and low exponent RSA vulnerabilities. *J. Cryptology 10*, 4 (1997), 233–260.

[15] DORNSEIF, M. 0wned by an iPod. Presentation, PacSec, 2004.

[16] DORNSEIF, M. Firewire – all your memory are belong to us. Presentation, CanSecWest/core05, May 2005.

[17] DWOSKIN, J., AND LEE, R. B. Hardware-rooted trust for secure key management and transient trust. In *Proc. 14th ACM Conference on Computer and Communications Security* (Oct. 2007), pp. 389–400.

[18] DYER, J. G., LINDEMANN, M., PEREZ, R., SAILER, R., VAN DOORN, L., SMITH, S. W., AND WEINGART, S. Building the IBM 4758 secure coprocessor. *Computer 34* (Oct. 2001), 57–66.

[19] ECKSTEIN, K., AND DORNSEIF, M. On the meaning of 'physical access' to a computing device: A vulnerability classification of mobile computing devices. Presentation, NATO C3A Workshop on Network-Enabled Warfare, Apr. 2005.

[20] ELSON, J., AND GIROD, L. Fusd – a Linux framework for user-space devices. http://www.circlemud.org/ jelson/software/fusd/.

[21] FERGUSON, N. AES-CBC + Elephant diffuser: A disk encryption algorithm for Windows Vista. http://www.microsoft.com/downloads/details.aspx?FamilyID=131dae03-39ae-48be-a8d6-8b0034c92555, Aug. 2006.

[22] GUTMANN, P. Secure deletion of data from magnetic and solid-state memory. In *Proc. 6th USENIX Security Symposium* (July 1996), pp. 77–90.

[23] GUTMANN, P. Data remanence in semiconductor devices. In *Proc. 10th USENIX Security Symposium* (Aug. 2001), pp. 39–54.

[24] IEEE 1619 SECURITY IN STORAGE WORKING GROUP. IEEE P1619/D19: Draft standard for cryptographic protection of data on block-oriented storage devices, July 2007.

[25] INTEL CORPORATION. Preboot Execution Enviroment (PXE) specification version 2.1, Sept. 1999.

[26] KENT, C. Draft proposal for tweakable narrow-block encryption. https://siswg.net/docs/LRW-AES-10-19-2004.pdf, 2004.

[27] LEE, R. B., KWAN, P. C., MCGREGOR, J. P., DWOSKIN, J., AND WANG, Z. Architecture for protecting critical secrets in microprocessors. In *Proc. Intl. Symposium on Computer Architecture* (2005), pp. 2–13.

[28] LIE, D., THEKKATH, C. A., MITCHELL, M., LINCOLN, P., BONEH, D., MITCHELL, J., AND HOROWITZ, M. Architectural support for copy and tamper resistant software. In *Symp. on Architectural Support for Programming Languages and Operating Systems* (2000), pp. 168–177.

[29] LINK, W., AND MAY, H. Eigenschaften von MOS-Ein-Transistorspeicherzellen bei tiefen Temperaturen. *Archiv für Elektronik und Übertragungstechnik 33* (June 1979), 229–235.

[30] LISKOV, M., RIVEST, R. L., AND WAGNER, D. Tweakable block ciphers. In *Advances in Cryptology – CRYPTO 2002* (2002), pp. 31–46.

[31] MACIVER, D. Penetration testing Windows Vista BitLocker drive encryption. Presentation, Hack In The Box, Sept. 2006.

[32] NATIONAL CONFERENCE OF STATE LEGISLATURES. State security breach notification laws. http://www.ncsl.org/programs/lis/cip/priv/breachlaws.htm, Jan. 2008.

[33] PETTERSSON, T. Cryptographic key recovery from Linux memory dumps. Presentation, Chaos Communication Camp, Aug. 2007.

[34] PTACEK, T. Recover a private key from process memory. http://www.matasano.com/log/178/recover-a-private-key-from-process-memory/.

[35] ROGAWAY, P. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In *Advances in Cryptology – ASIACRYPT 2004* (2004), pp. 16–31.

[36] RSA LABORATORIES. PKCS #1 v2.1: RSA cryptography standard. ftp://ftp.rsasecurity.com/pub/pkcs/pkcs-1/pkcs-1v2-1.pdf.

[37] SCHEICK, L. Z., GUERTIN, S. M., AND SWIFT, G. M. Analysis of radiation effects on individual DRAM cells. *IEEE Transactions on Nuclear Science 47* (Dec. 2000), 2534–2538.

[38] SEAGATE CORPORATION. Drivetrust technology: A technical overview. http://www.seagate.com/docs/pdf/whitepaper/TP564_DriveTrust_Oct06.pdf.

[39] SHAMIR, A., AND VAN SOMEREN, N. Playing "hide and seek" with stored keys. *Lecture Notes in Computer Science 1648* (1999), 118–124.

[40] SKOROBOGATOV, S. Low-temperature data remanence in static RAM. University of Cambridge Computer Laboratory Technical Report No. 536, June 2002.

[41] SMITH, S. W. *Trusted Computing Platforms: Design and Applications*, first ed. Springer, 2005.

[42] TRUSTED COMPUTING GROUP. Trusted Platform Module specification version 1.2, July 2007.

[43] VIDAS, T. The acquisition and analysis of random access memory. *Journal of Digital Forensic Practice 1* (Dec. 2006), 315–323.

[44] WEINMANN, R.-P., AND APPELBAUM, J. Unlocking FileVault. Presentation, 23rd Chaos Communication Congress, Dec. 2006.

[45] WEINMANN, R.-P., AND APPELBAUM, J. VileFault. http://vilefault.googlecode.com/, Jan. 2008.

[46] WYNS, P., AND ANDERSON, R. L. Low-temperature operation of silicon dynamic random-access memories. *IEEE Transactions on Electron Devices 36* (Aug. 1989), 1423–1428.

An Introduction to New Stream Cipher Designs

Tor E. Bjørstad

University of Bergen, Norway
Email : tor.bjorstad@ii.uib.no

1 What is a stream cipher?

Even with "nothing to hide", it is very often desirable to protect the privacy of our data and our communications. The usual way to do this is by using encryption[1]. It is safe to say that the Internet as we know it would not exist without strong crypto – whether it is used to protect remote logins, e-commerce transactions, the hard disk of your laptop, or something completely different. The goal is mostly the same in all cases: keeping the data or communications confidential.

Briefly speaking, a symmetric encryption algorithm takes some *plaintext* data and a secret *key* as input, and outputs a *ciphertext*. The goal of the exercise, of course, is that anyone who does not know the key, should not be able to deduce anything useful about the plaintext or the key from the ciphertext, while anyone who does know the key is able to decrypt and recover the plaintext from the ciphertext. There are two main classes of symmetric encryption algorithms, block ciphers and stream ciphers.

Block ciphers are the most well known type of symmetric encryption. Some famous algorithms are the old Data Encryption Standard (DES) and its replacement, the Advanced Encryption Standard (AES). These ciphers operate, as the name indicates, on fixed-length data blocks. AES takes 128 bits of plaintext as input together with a 128-bit secret key, and outputs 128 bits of ciphertext. In order to encrypt larger amounts of data securely, a chaining mode such as CBC must be used. Data fragments shorter than the block length usually have to be padded to the block length, increasing transmission overhead.

Stream ciphers are different. A stream cipher may, very loosely speaking, be thought of as a cryptographically secure pseudo-random number generator with some extra bells and whistles. These algorithms take the secret key as well as a public initialisation vector (IV, sometimes called a *nonce*) as input[2], and output a stream of random-looking symbols, known as the *keystream*. The

[1] It should be emphasised at this point that encryption alone is almost *never* enough by itself to build a secure protocol or application. Generally, one should *always* use a digital signature or a message authentication code (MAC) to maintain the integrity of encrypted data. Surprisingly often, it will be possible to attack an encrypted protocol without integrity checking, by tampering with the data in some particular way. However, this is a completely different cup of tea and far outside the scope of this short paper.

[2] The use of an IV makes it possible to encrypt several data streams without changing long-term secret keys. However, a given key/IV pair must only be used once.

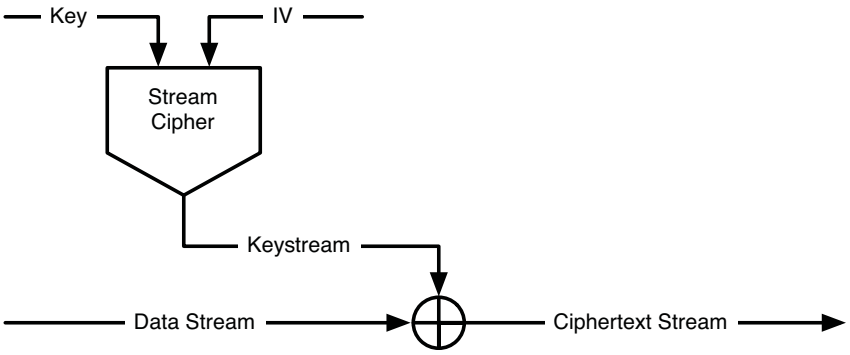

Fig. 1. Schematic representation of encryption with a stream cipher.

keystream symbols are usually either single bits, bytes, or machine words. To encrypt a data stream, one simply has to exclusive-or (XOR) the data symbols with the keystream. Decryption is of course the exact opposite, since the XOR operation is symmetric; when the ciphertext c is computed from the message m and keystream x as $c = m \bigoplus x$, we also have that $m = c \bigoplus x$. Real-world examples of stream ciphers include RC4 (used for WEP/WPA, by Bittorrent, and by SSL, to name a few) A5/1 (GSM telephony) and E0 (Bluetooth), as well as AES in some modes of operation (notably counter, or CTR-mode).

How does a stream cipher work? Although the specific details of stream ciphers vary immensely fro cipher to cipher, there is a certain amount of common structure. A stream cipher consists of a certain amount of internal state, which should be at least twice the size of the secret key to prevent certain generic attacks. Given a key and IV, the algorithm proceeds by a specified number of initialization steps, in which the key, IV and initial contents of the state are mixed in a nonlinear fashion. After this, the cipher starts outputting keystream symbols as a function of the (now sufficiently randomised) state, while continuing to mix and evolve the contents of the state itself. A modern stream cipher specification should make clear certain usage limits: how many keystream bits can be generated by a single key/IV pair, and how many IVs can be used before the secret key itself must be changed.

What are the advantages of stream ciphers over block ciphers? Stream ciphers may be faster or have a smaller implementation footprint than comparable block ciphers. They operate more naturally on data of short, variable or unknown length. Finally, the keystream generation is completely independent of the plaintext data, and so it may be computed in parallel with or in advance of the data stream. In general, it is also useful for system designers to have a reasonable selection of different encryption algorithms to choose between, as this makes it possible both to select precise performance tradeoffs suitable for a specific application, and to avoid a cryptographic monoculture where everyone is

using AES and subsequently everyone gets in big trouble if future cryptanalysis reveals weaknesses in AES.

2 What is a secure stream cipher?

The usual starting assumption for attacks on stream cipher, is that the attacker has access to large amounts of keystream, generated under a number of different IVs which may (if necessary) be adaptively chosen by the attacker. In a sense this is a very generous setting; it is to be hoped that an actual real-world attack scenario will be (much) less bad. But it corresponds well with Kerckhoff's second principle: one always should assume that the enemy knows everything about the encryption system used, apart from the secret key itself. Conversely, if there are no attacks that can be applied even in this rather theoretical setting, there will surely not be any attacks in a more restricted (and possibly more realistic) situation.

There are two main criteria for the security of a stream cipher with a k-bit key. Firstly, the attacker should not be able to predict future keystream output by the cipher, whether this happens by recovering the secret key, recovering the internal state of the cipher at some point, or otherwise. The attacker can obviously do this by testing all possible secret keys, so the complexity of a brute force attack (requiring at most 2^k executions of the algorithm) gives a performance baseline to which any alleged attack should be compared.

Secondly, the attacker should not be able to distinguish keystream from random under the given usage limits for the cipher. While a distinguishing attack is certainly less serious than a full state recovery, it does indicate that the algorithm has some kind of undesirable internal structure. It should be emphasised that the existence of an "attack" of any kind on a stream cipher, does not mean that the attack is *practical* in any way. Rather, it implies that the algorithm is strictly worse than a hypothetical *ideal* stream cipher, where the only applicable attacks are generic attacks such as brute-forcing the key, and the keystream is truly indistinguishable from random when the key is unknown.

3 What is the current state of the art?

The most widely used stream cipher around is, by far, RC4. Designed by Ronald Rivest in 1987, it is extremely fast in software and can be implemented in just a few lines of code. The history of RC4 is actually quite interesting, as the cipher was originally considered a trade secret by RSA, and only became public knowledge after it was anonymously leaked on the Internet in 1994 [14]. Unfortunately, the age of RC4 is increasingly starting to show. For one, the cipher specification does not specify how to use an initialisation vector with the algorithm, which means that implementors have to be very careful about how they do this if they want to generate multiple keystreams based on the same long-term secret. More seriously, it was discovered in 2001 that the RC4 keystream exhibits various statistical biases that can be used to distinguish it form random and relate it to

the underlying secret key [10], and improved attacks along these lines have also been found since then. A usual technique to mitigate these attacks is to discard the first N bytes of the RC4 keystream; a typical value is $N = 1024$.

These issues imply that RC4 can not be considered (theoretically) secure by modern standards, as discussed in the previous section. Even though it may still be possible to use RC4 in a sufficiently secure way by a careful (and lucky) implementor, the WEP fiasco as well as the recent attack on WPA by Beck and Tews [2] show that the theoretical weaknesses of RC4 also lead to practical attacks on protocols in which RC4 is used. From an academic point of view, certainly, RC4 should not be used for new applications – even though it offers very attractive performance and ease of implementation from an engineering point of view, it is simply considered too risky and too difficult to get right.

Unfortunately, the state of other popular stream ciphers is no less dire. The A5/1 cipher and its variants used in GSM have a completely inadequate key length of 54 bits, and additional attacks have been found which are faster than brute force. Similarly, E0 has been broken by cryptanalysts. The European Union-based NESSIE project [12], which was aimed at evaluating the security of various cryptographic primitives, did not recommend any stream ciphers in their final report, because all the submitted algorithms were successfully attacked. Although there do exist various other stream ciphers around that are still unbroken (notably SNOW 2.0 [13], a tweaked variant of one of the NESSIE ciphers), none of them have really gained widespread acceptance and recognition outside the academic community.

The only popular, secure and widespread "stream cipher" that remains is, in the author's opinion, AES (or any other secure block cipher, but AES is after all the standard) operating in counter (CTR) mode. But this yields the obvious question: is it really not possible to design a secure, special-purpose stream cipher, which is more efficient than what you get by adapting a block cipher to the task?

4 What is the eSTREAM project?

The eSTREAM project [9] was launched as part of the EU-funded network of excellence ECRYPT [8], which ran from 2004 to 2008. This was partly in response to the dismal showing of the NESSIE stream ciphers, and had a stated project goal to identify "new stream ciphers that might become suitable for widespread adoption".

A call for primitives was put forward in the fall of 2004, and attracted 34 submissions from all over the world. The candidate algorithms were divided in two categories, or profiles, one for software-oriented algorithms, and one for ciphers suitable for hardware implementation. The design goals for the two profiles were somewhat different. Software candidates should offer a key size of at least 128 bits, and provide some significant advantage over the state of the art (i.e. AES-CTR) with respect to throughput. Candidates for the hardware profile should outperform AES in restricted environments with respect to relevant parameters

such as gate count, power consumption and speed, and provide a key size of at least 80 bits.

After three evaluation phases, the eSTREAM project ended in the spring of 2008 with a final portfolio of 8 "promising" ciphers, 4 in each profile. One of these, F-FCSR, was removed from the portfolio in September 2008 after new cryptanalytic results were found. It is important to remember that, unlike the previous AES competition and the current SHA-3 competition, the goal of eSTREAM was not to develop a new international standard for stream ciphers, but merely to act as a focus for academic interest, and attempt to identify the best candidates among the various designs. While the different eSTREAM algorithms are still quite new and untested and new weaknesses may yet be found, the portfolio can be considered to represent the current state of academic research on stream ciphers.

The eSTREAM portfolio consists, as of November 2008, of the following seven stream ciphers:

- HC-128 [15] (software), supporting 128-bit keys.
- Rabbit [6] (software), supporting 128-bit keys.
- Salsa20/12 [5] (software), supporting 128 and 256-bit keys.
- SOSEMANUK [3] (software), supporting 128-256 bit keys.
- Trivium [7] (hardware), supporting 80-bit keys.
- Grain v1 [11] (hardware), supporting 80-bit keys.
- MICKEY v2 [1] (hardware), supporting 80-bit keys.

In the accompanying lecture, the aim of the author is to give a lightning tour of the eSTREAM portfolio ciphers, emphasising their respective strengths and weaknesses, and examining how the different algorithms are constructed.

References

1. S. Babbage and M. Dodd. The MICKEY stream ciphers. *Lecture Notes in Computer Science*, 4986:191–209, 2008. `http://www.ecrypt.eu.org/stream/mickeypf.html`.
2. M. Beck and E. Tews. Practical attacks against WEP and WPA, 2008. `http://dl.aircrack-ng.org/breakingwepandwpa.pdf`.
3. C. Berbain, O. Billet, A. Canteaut, N. Courtois, H. Gilbert, L. Goubin, A. Gouget, L. Granboulan, C. Lauradoux, M. Minier, T. Pornin, and H. Sibert. SOSEMANUK, a fast software-oriented stream cipher. *Lecture Notes in Computer Science*, 4986:98–118, 2008. `http://www.ecrypt.eu.org/stream/sosemanukpf.html`.
4. D. Bernstein. Notes on the ECRYPT Stream Cipher project (eSTREAM). `http://cr.yp.to/streamciphers.html`.
5. D. Bernstein. The Salsa20 family of stream ciphers. *Lecture Notes in Computer Science*, 4986:84–97, 2008. `http://www.ecrypt.eu.org/stream/salsa20pf.html`.
6. M. Boesgaard, M. Vesterager, and E. Zenner. The Rabbit stream cipher. *Lecture Notes in Computer Science*, 4986:69–83, 2008. `http://www.ecrypt.eu.org/stream/rabbitpf.html`.
7. C. de Cannière and B. Preneel. TRIVIUM. *Lecture Notes in Computer Science*, 4986:244–266, 2008. `http://www.ecrypt.eu.org/stream/triviumpf.html`.

8. ECRYPT Network of Excellence in Cryptology. http://www.ecrypt.eu.org/.
9. The eSTREAM project. http://www.ecrypt.eu.org/stream/.
10. S. Fluhrer, I. Mantin, and V. Shoup. Weaknesses in the key scheduling algorithm of RC4. In *Proceedings of SAC 2001*, volume 2259 of *Lecture Notes in Computer Science*, pages 1–24. Springer–Verlag, 2001.
11. M. Hell, T. Johansson, A. Maximov, and W. Meier. The Grain family of stream ciphers. *Lecture Notes in Computer Science*, 4986:179–190, 2008. http://www.ecrypt.eu.org/stream/grainpf.html.
12. New European Schemes for Signatures, Integrity and Encryption. http://www.cosic.esat.kuleuven.be/nessie/.
13. T. Johansson P. Ekdahl. A new version of the stream cipher SNOW. In *Proceedings of SAC 2002*, volume 2595 of *Lecture Notes in Computer Science*, pages 47–61. Springer–Verlag, 2002. http://www.it.lth.se/cryptology/snow/.
14. RC4 algorithm revealed. http://groups.google.com/group/sci.crypt/msg/10a300c9d21afca0.
15. H. Wu. The stream cipher HC-128. *Lecture Notes in Computer Science*, 4986:39–47, 2008. http://www.ecrypt.eu.org/stream/hcpf.html.

Reverse-Engineering a Cryptographic RFID Tag

Karsten Nohl and David Evans
University of Virginia
Department of Computer Science
{nohl,evans}@cs.virginia.edu

Starbug and Henryk Plötz
Chaos Computer Club
Berlin
starbug@ccc.de, henryk@ploetzli.ch

Abstract

The security of embedded devices often relies on the secrecy of proprietary cryptographic algorithms. These algorithms and their weaknesses are frequently disclosed through reverse-engineering software, but it is commonly thought to be too expensive to reconstruct designs from a hardware implementation alone. This paper challenges that belief by presenting an approach to reverse-engineering a cipher from a silicon implementation. Using this mostly automated approach, we reveal a cipher from an RFID tag that is not known to have a software or micro-code implementation. We reconstruct the cipher from the widely used Mifare Classic RFID tag by using a combination of image analysis of circuits and protocol analysis. Our analysis reveals that the security of the tag is even below the level that its 48-bit key length suggests due to a number of design flaws. Weak random numbers and a weakness in the authentication protocol allow for pre-computed rainbow tables to be used to find any key in a matter of seconds. Our approach of deducing functionality from circuit images is mostly automated, hence it is also feasible for large chips. The assumption that algorithms can be kept secret should therefore to be avoided for any type of silicon chip.

Il faut qu'il n'exige pas le secret, et qu'il puisse sans inconvénient tomber entre les mains de l'ennemi.
([A cipher] must not depend on secrecy, and it must not matter if it falls into enemy hands.)
August Kerckhoffs, *La Cryptographie Militaire*, January 1883 [13]

1 Introduction

It has long been recognized that security-through-obscurity does not work. However, vendors continue to believe that if an encryption algorithm is released only as a hardware implementation, then reverse-engineering the cipher from hardware alone is beyond the capabilities of likely adversaries with limited funding and time. The design of the cipher analyzed in this paper, for example, had not been disclosed for 14 years despite more than a billion shipped units. We demonstrate that the cost of reverse engineering a cipher from a silicon implementation is far lower than previously thought.

In some cases, details of an unknown cryptographic cipher may be found by analyzing the inputs and outputs of a black-box implementation. Notable examples include Bletchley Park's breaking the Lorenz cipher during World War II without ever acquiring a cipher machine [5] and the disclosure of the DST cipher used in cryptographic Radio Frequency Identification (RFID) tokens from Texas Instruments [4]. In both cases, researchers started with a rough understanding of the cipher's structure and were able to fill in the missing details through cryptanalysis of the cipher output for known keys and inputs. This black-box approach requires some prior understanding of the structure of a cipher and is only applicable to ciphers with statistical weaknesses. The output of a sound cipher should not be statistically biased and therefore should not leak information about its structure.

Other ciphers have been disclosed through disassembly of their software implementation. Such implementations can either be found in computer software or as microcode on an embedded micro-controller. Ciphers found through software disassembly include the A5/1 and A5/2 algorithms that secure GSM cell phone communication [1] and the Hitag2 and Keeloq algorithms used in car remote controls [3]. The cryptography on the RFID tags we analyzed is not known to be available in software or in a micro-code implementation; tags and reader chips implement the cipher entirely in hardware.

In this paper, we focus on revealing proprietary cryptography from its silicon implementation alone. Reverse-engineering silicon is possible even when very little is known about a cipher and no software implementation

exists. The idea of reverse-engineering hardware is not new. Hardware analysis is frequently applied in industry, government, and the military for spying, security assessments, and protection of intellectual property. Such reverse-engineering, however, is usually considered prohibitively expensive for typical attackers, because of the high prices charged by professionals offering this service. The key contribution of this work is demonstrating that reverse-engineering silicon is cheap and that it can be mostly automated. This is the first published work to describe the details of reverse-engineering a cryptographic function from its silicon implementation. We describe a mostly automated process that can be used to cheaply determine the functionality of previously unknown cipher implementations.

We demonstrate the feasibility of our approach by revealing the cipher implemented on the NXP Mifare Classic RFID tags, the world's most widely used cryptographic RFID tag [16]. Section 2 describes our reverse-engineering method and presents the cipher. Section 3 discusses several weaknesses in the cipher beyond its short key size. Weak random numbers combined with a protocol flaw allow for rainbow tables to be computed that reduce the attack time from weeks to minutes. Section 4 discusses some potential improvements and defenses. While we identify fixes that would increase the security of the Mifare cipher significantly, we conclude that good security may be hard to achieve within the desired resource constraints.

2 Mifare Crypto-1 Cipher

We analyzed the Mifare Classic RFID tag by NXP (formerly Philips). This tag has been on the market for over a decade with over a billion units sold. The Mifare Classic card is frequently found in access control systems and tickets for public transport. Large deployments include the Oyster card in London, and the SmartRider card in Australia. Before this work, the Netherlands were planning to deploy Mifare tags in OV-chipkaart, a nationwide ticketing system, but the system will likely be reengineered after first news about a potential disclosure of the card's details surfaced [17]. The Mifare Classic chip currently sells for 0.5 Euro in small quantities, while tags with larger keys and established ciphers such as 3-DES are at least twice as expensive.

The cryptography found in the Mifare cards is a stream cipher with 48-bit symmetric keys. This key length has been considered insecure for some time (for example, the Electronic Frontier Foundation's DES cracking machine demonstrated back in 1998 that a moderately-funded attacker could brute force 56-bit DES [6]) and the practical security that Mifare cards have experienced in the past relies primarily on the belief that its cipher was secret. We find that the security of the Mifare Classic is even weaker than the short key length suggests due to flaws in its random number generation and the initialization protocol discussed in Section 3.

The data on the Mifare cards is divided into sectors, each of which holds two different keys that may have different access rights (e.g., read/write or read-only). This division allows for different applications to each store encrypted data on a tag—an option rarely used in practice. All secrets are set to default values at manufacturing time but changed before issuing the tags to users. Different tags in a system may share the same read key or have different keys. Sharing read keys minimizes the overhead of key-distribution to offline readers. We find, however, that the protocol level measures meant to prevent different users from impersonating each other are insufficient. Unique read and write keys should, therefore, be used for each tag and offline readers should be avoided as much as possible.

2.1 Hardware Analysis

The chip on the Mifare Classic tag is very small with a total area of roughly one square millimeter. About a quarter of the area is used for 1K of flash memory (a 4K version is also available); another quarter is occupied by the radio front-end and outside connectivity, leaving about half the chip area for digital logic including cryptography.

The cryptography functions make up about 400 2-NAND gate equivalents (GE), which is very small even compared to highly optimized implementations of standard cryptography. For example, the smallest known implementation of the AES block cipher (which was specifically designed for RFID tags) requires 3400 GEs [7]. The cryptography on the Mifare tags is also very fast and outputs 1 bit of key stream in every clock cycle. The AES circuit, by comparison, takes 1000 clock cycles for one 128-bit AES operation (10 milliseconds on a tag running at 106 kHz).

To reverse engineer the cryptography, we first had to get access to sample chips, which are usually embedded in credit card size plastic cards. We used acetone to dissolve the plastic card, leaving only the blank chips. Acetone is easier and safer to handle than alternatives such as fuming nitric acid, but still dissolves plastic cards in

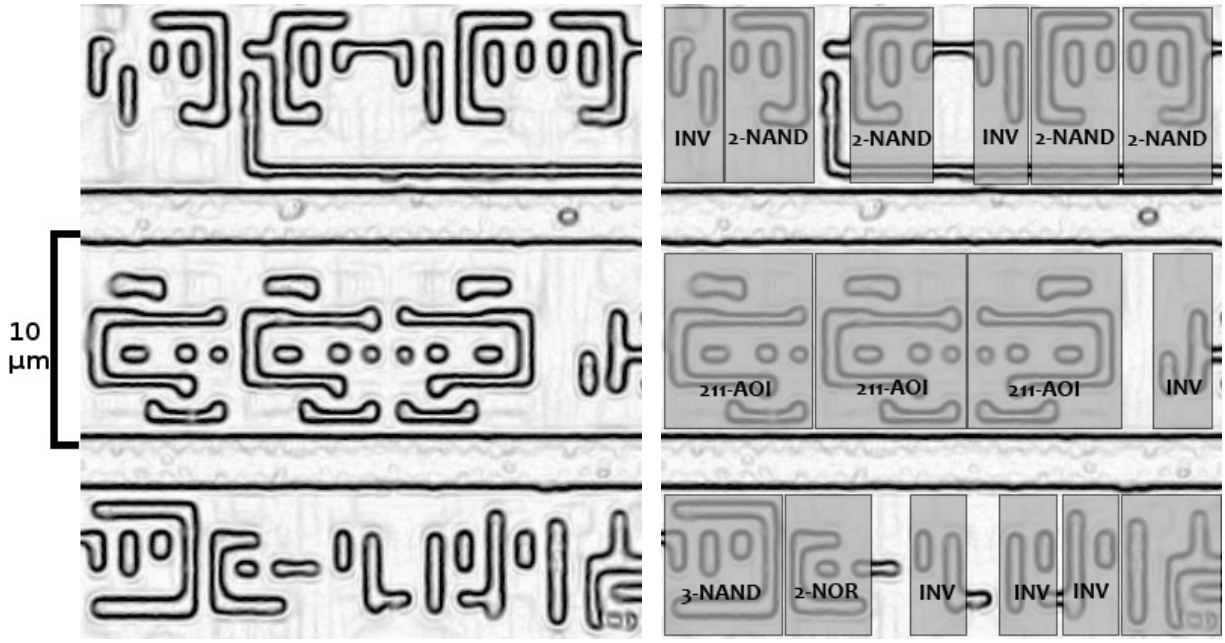

(a) (b)

Figure 1: (a) Source image of layer 2 after edge detection; (b) after automated template detection.

about half an hour. Once we had isolated the silicon chips, we removed each successive layer through mechanical polishing, which we found easier to control than chemical etching. Simple polishing emulsion or sandpaper with very fine grading of $0.04\mu m$ suffices to take off micrometer-thick layers within minutes.

Although the polishing is mostly straightforward, the one obstacle to overcome is the chip tilting. Since the chip layers are very close together, even the smallest tilt leads to cuts through several layers. We addressed this problem in two ways. First, we embedded the millimeter-size chip in a block of plastic so it was easier to handle. Second, we accpeted that we could not completely avoid tilt using our simple equipment and adapted our image stitching tools to patch together chip layers from several sets of pictures, each imaging parts of several layers.

The chip contains a total of six layers, the lowest of which holds the transistors. We took pictures using a standard optical microscope at a magnification of 500x. From multiple sets of these images we were able to automatically generate images of each layer using techniques for image tiling that we borrowed from panorama photography. We achieved the best results using the open source tool `hugin` (http://hugin.sourceforge.net/) by setting the maximum variance in viewer angle to a very small value (e.g., 0.1°) and manually setting a few control points on each image.

The transistors are grouped in gates that each perform a logic function such as AND, XOR, or flip-flop as illustrated in Figure 1. Across the chip there are several thousand such logic gates, but only about 70 different types of gates. As a first step toward reconstructing the circuit, we built a library of these gates. We implemented template matching that given one instance of a logic gate finds all the other instances of the same gate across the chip. Our tools take as input an image of layer 2, which represents the logic level, and the position of instances of different logic gates in the image. The tools then use template matching to find all other instances of the gate across the image, including rotated and mirrored variants. Since larger gates sometimes contain smaller gates as building blocks, the matching is done in order of decreasing gate sizes.

Our template matching is based on *normalized cross-correlation* which is a well-known similarity test [14] and implemented using the MATLAB image processing library. Computing this metric is computationally more complex than standard cross-correlation, but the total running time of our template matching is still under ten minutes for the whole chip. Normalized cross-correlation is insensitive to the varying brightness across our different images and the template matching is able to find matches with high accuracy despite varying coloration and distortion of the structures that were caused by the polishing.

We then manually annotated each type of gate with its respective functionality. This step could be automated as well through converting the silicon-level depiction of each gate into a format suitable for a circuit simulation program. We decided against this approach because the overhead seemed excessive. For larger libraries that perhaps intentionally vary the library cells in an attempt to impede reverse-engineering, however, automation is certainly possible and has already been demonstrated in other projects [2].

Our template matching provides a map of the different logic gates across the chip. While it would certainly have been possible to reverse-engineer the whole RFID tag, we focused our attention on finding and reconstructing the cryptographic components. We knew that the stream cipher would have to include at least a 48-bit register and a number of XOR gates. We found these components in one of the corners of the chip along with a circuit that appeared to be a random number generator as it has an output, but no input.

Focusing our efforts on only these two parts of the chip, we reconstructed the connections between all the logic gates. This step involved considerable manual effort and was fairly error-prone. All the errors we made were found through a combination of redundant checking and statistical tests for some properties that we expected the cipher to have such as an even output distribution of blocks in the filter function. We have since implemented scripts to automate the detection of wires, which can speed the process and improve its accuracy. Using our manually found connections as ground truth we find that our automated scripts detect the metal connection and intra-layer vias correctly with reasonably high probability. In our current tests, our scripts detect over 95% of the metal connections correctly and the few errors they make were easily spotted manually by overlaying the source image and the detection result. These results are, however, preliminary, as many factors are not yet accounted for. To assess the potential for automation more thoroughly, we plan to test our tools on different chips, using different imaging systems, and having different users check the results.

In the process of reconstructing the circuit, we did not encounter any added obscurity or tamper-proofing. Because the cryptographic components are highly structured, they were particularly easy to reconstruct. Furthermore, we could test the validity of different building blocks by checking certain statistical properties. For example, the different parts of the filter function each have an even output distribution so that the output bits are not directly disclosing information about single state bits.

The map of logic gates and the connections between them provides us with almost enough information to discover the cryptographic algorithm. Because we did not reverse-engineer the control logic, we do not know the exact timing and inputs to the cipher. Instead of reconstructing more circuitry, we derived these missing pieces of information from protocol layer communication between the Mifare card and reader.

2.2 Protocol Analysis

From the discovered hardware circuit, we could not derive which inputs are shifted into the cipher in what order, partly because we did not reverse the control logic, but also because even with complete knowledge of the hardware we would not yet have known what data different memory cells contain. To add the missing details to the cipher under consideration, and to verify the results of the hardware analysis, we examined communication between the Mifare tags and a Mifare reader chip.

An NXP reader chip is included on the OpenPCD open source RFID reader, whose flexibility proved to be crucial for the success of our project. The OpenPCD includes an ARM micro-controller that controls the communication between the NXP chip and the Mifare card. This setup allows us to record the communication and provides full control over the timing of the protocol. Through timing control we can amplify some of the vulnerabilities we discovered as discussed in Section 3.

No details of the cipher have been published by the manufacturer or had otherwise been leaked to the public prior to this work. We guessed that the secret key and the tag ID were shifted into the shift register sequentially rather than being combined in a more complicated way. To test this hypothesis, we checked whether a reader could successfully authenticate against a tag using an altered key and an altered ID. Starting with single bit changes in ID and key and progressively extending our search to larger variations, we found a number of such combinations that indeed successfully authenticated the reader to the tag. From the pattern of these combinations we could derive not just the order of inputs, but also the structure of the linear feedback shift register, which we had independently found on the circuit level. Combining these insights into the authentication protocol with the results of our hardware analysis gave us the whole Crypto-1 stream cipher, shown in Figure 2.
The cipher is a single 48-bit linear feedback shift register (LFSR). From a fixed set of 20 state bits, the one bit of key stream is computed in every clock cycle. The shift register has 18 taps (shown as four downward arrows in

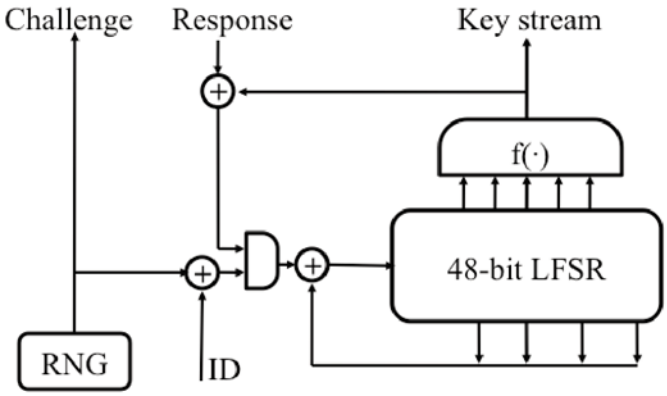

Figure 2: Crypto-1 stream cipher and initialization.

the figure) that are linearly combined to fill the first register bit on each shift. The update function does not contain any non-linearity, which by today's understanding of cipher design can be considered a serious weakness. The generating polynomial of the register is (with x^i being the ith bit of the shift register):

$$\begin{aligned} & x^{48}+x^{43}+x^{39}+x^{38}+x^{36}+x^{34}+x^{33}+x^{31}+x^{29} \\ + \quad & x^{24}+x^{23}+x^{21}+x^{19}+x^{13}+x^{9}+x^{7}+x^{6}+x^{5}+1. \end{aligned}$$

The polynomial is *primitive* in the sense that it is irreducible and generates all $\left(2^{48}-1\right)$ possible outputs in succession. To confirm this, we converted the Fibonacci LFSR into a Galois LFSR for which we can compute any number of steps in a few Galois field multiplications. We then found that the cipher state repeats after $\left(2^{48}-1\right)$ steps, but not after any of the possible factors for this number. The LSFR is hence of maximum-length.

The protocol between the Mifare chip and reader loosely follows the ISO 9798-2 specification, which describes an abstract challenge-response protocol for mutual authentication. The authentication protocol takes a shared secret key and a unique tag ID as its inputs. At the end of the authentication, the parties have established a session key for the stream cipher and both parties are convinced that the other party knows the secret key.

3 Cipher Vulnerabilities

The 48-bit key used in Mifare cards makes brute-force key searches feasible. Cheaper than brute-force attacks, however, are possible because of the cipher's weak cryptographic structure. While the vulnerability to brute-force attacks already makes the cipher weak, the cheaper attacks are relevant for many Mifare deployments such as fare collection where the value of breaking a particular key is relatively low. Weaknesses of the random number generator and the cryptographic protocol allow an attacker to pre-compute a codebook and perform key-lookups quickly and cheaply using rainbow tables.

3.1 Brute-Force Attack

In a brute-force attack an attacker records two challenge-response exchanges between the legitimate reader and a card and then tries all possible keys for whether they produce the same result.

To estimate the expected time for a brute-force attack, we implemented the cipher on FPGA devices by Pico Computing. Due to the simplicity of the cipher, 6 fully-pipelined instances can be squeezed into a single Xilinx Virtex-5 LX50 FPGA. Running the implementation on an array of 64 such FPGAs to try all 2^{48} keys takes under 50 minutes.

3.2 Random Number Generation

The random number generator (RNG) used on the Mifare Classic tags is highly insecure for cryptographic applications and further decreases the attack complexity by allowing an attacker to pre-compute a codebook.
The random numbers on Mifare Classic tags are generated using a linear feedback shift register with constant initial condition. Each random value, therefore, only depends on the number of clock cycles elapsed between the time the tag is powered up (and the register starts shifting) and the time the random number is extracted. The numbers are generated using a maximum length 16-bit LFSR of the form:

$$x^{16}+x^{14}+x^{13}+x^{11}+1.$$

The register is clocked at 106 kHz and wraps around every 0.6 seconds after generating all 65,535 possible output values. Aside from the highly insufficient length of the random numbers, an attacker that controls the timing of the protocol controls the generated number. The weakness of the RNG is amplified by the fact that the generating LFSR is reset to a known state every time the tag starts operating. This reset is completely unnecessary, involves hardware overhead, and destroys the randomness that previous transactions and unpredictable noise left in the register.

We were able to control the number the Mifare random number circuit generated using the OpenPCD reader and custom-built firmware. In particular, we were able to generate the same "random" nonce in each query, thereby completely eliminating the tag randomness from the authentication process. Moreover, we found the same weakness in the 32-bit random numbers generated by the reader chip, which suggests that a similar hardware implementation is used in the chip and reader. Here, too, we were able to repeatedly generate the same number. While in our experiments this meant controlling the timing of the reader chip, a skilled attacker will likely be able to exploit this vulnerability even in realistic scenarios where no such control over the reader is given. The attacker can predict forthcoming numbers from the numbers already seen and precisely chose the time to start interacting with the reader in order to receive a certain challenge. The lack of true randomness on both reader and tag enable an attacker to eliminate any form of randomness from the authentication protocol. Depending on the number of precomputed codebooks, this process might take several hours and the attack might not be feasible against all reader chips.

3.3 Pre-Computing Keys

Several weaknesses of the Mifare card design add up to what amounts to a full codebook pre-computation. First, the key space is small enough for all possible keys to be included. Second, the random numbers are controllable. In addition, the secret key and the tag ID are combined in such a way that for each session key there exists exactly one key for each ID that would result in that session key. The key and the ID are shifted into the register sequentially, but no non-linearity is mixed in during this process. As explained in Section 2.2, for every delta of ID bits, there exists a delta of key bits that corrects for the difference and results in the same session key. Therefore, given a key that for some ID results in a session key, there exists a key for any ID that would result in the same session key. This bijective mapping allows for a codebook that was pre-computed for only a single ID to be used to find keys for all other IDs as well.

A codebook for all keys would occupy 1500 Terabytes, but can be stored more economically in rainbow tables. Rainbow tables store just enough information from a key space for finding any key with high probability, but require much less space than a table for all keys [9, 15]. Each "rainbow" in these tables is the repeated application of slight variants of a cryptographic operation. In our case, we start with a random key and generate the output of the authentication protocol for this key, then use this output as the next key for the authentication, generate its output, use that as the next key, and so on. We then only store the first and last value of each rainbow, but compute enough rainbows so that almost all keys appear in one of them. To find a key from such a rainbow table, a new rainbow is computed starting at a recorded output from the authentication protocol. If any one of the generated values in this series is also found in the stored end values of the rainbows, then the key used in the authentication protocol can be found from the corresponding start values of that matching rainbow. The time needed to find a key grows as the size of the tables shrinks.

Determining any card's secret key will be significantly cheaper than trying out all possible keys even for rainbow tables that only occupy a few Terabytes and can be almost as cheap as a database lookup. The fact that an attacker can use a pre-computed codebook to reveal the keys from many cards dramatically changes the economics of an attack in favor of the attacker. This means that even attacks on low-value cards like bus tickets might be profitable.

3.4 Threat Summary

To summarize the threat to systems that rely on Mifare encryption for security, we illustrate a possible attack. An attacker would first scan the ID from a valid card. This number is unprotected and always sent in the clear. Next, the attacker would pretend to be that card to a legitimate reader, record the reader message of the challenge-response protocol with controlled random nonces, and abort the transaction. Given only two of these messages, the key of the card can be found in the pre-computed rainbow tables in a matter of minutes and then used to read the data from the card. This gives the attacker all the information needed to clone the card.

4 Discussion

The illustrated attack is yet another example of security-by-obscurity failing. Weaknesses in the exposed cipher reveal the pitfalls of proprietary cipher design without peer-review. A few changes in the design would have made some of the discussed attacks infeasible and could have increased the key size within the same hardware constraints to make brute-force attacks less likely. Much better security, however, can only be achieved through better, more thoroughly analyzed ciphers.

4.1 Potential Fixes

The system is vulnerable against codebook attacks because of its weak random numbers and the linear combination of key and ID. Both can be fixed without adding extra hardware or slowing down the operation.

Better, yet still not cryptographically sound, random numbers can be generated by exploiting the fact that memory cells are initially in an undetermined state [10]. The same behavior can be caused in flip-flops like those that make up the state register of the stream cipher simply by not resetting the flip-flops at initialization time. The cipher state would start in a random state and then evolve using the cipher's feedback loop until a random number is needed. At this point, the register contains a mostly unpredictable number of the size of the state register.

Because this design generates random numbers within the same registers that are used for the cipher states, it eliminates the need for a separate additional PRNG circuit. The saved area could then be spent on increasing the size of the cipher state. In the area of the 48-bit Crypto-1 and its 16-bit RNG, a 64-bit stream cipher that also produces significantly better pseudo-random number could hence be implemented. This increases the size and quality of the random numbers and at the same time increases the key size beyond the point where brute-force attacks can be done cheaply.

To further improve the resistance against codebook attacks, the non-linear feedback should be combined with either key or ID when shifted into the register to break the bijective mapping between different key-ID pairs. This measure does not increase implementation costs, since we only integrate the output of the filter function which is already computed.

To improve the resistance of the cipher against statistical attacks, the update function must be made non-linear, either by feeding some intermediate result of the filter function into the linear register or by using a non-linear feedback shift register instead.

None of the possible fixes will make the cipher appropriate for high security applications, but they improve the resistance against the most concerning attacks and can be done without any additional implementation cost.

4.2 Possible Defenses

Possible ways to protect against the described attacks include using standard, peer-reviewed, established cryptography such as the 3-DES block cipher that is already found on some of the more expensive cards including some of the Mifare line of products. A cheaper alternative that can be implemented in about twice the size of Crypto-1 is the Tiny Encryption Algorithm (TEA) [12, 18]. This established low-cost block cipher has publicly been scrutinized for several years and is so far only known to be vulnerable to some expensive attacks [11]. While TEA is far more secure than Crypto-1, it is also much slower. A Mifare authentication takes little more than one millisecond, while a minimum-size implementation of TEA would take about ten times as long. This would still be fast enough for most applications where Mifare cards are currently used.

Other known ways to protect against card cloning include fraud detection algorithms that are widely used in monitoring credit card transactions. These algorithms detect unusual behavior and can prevent fraudulent transactions, but require storing and analyzing transaction data, which runs contrary to the desire for privacy in RFID applications. Fraud protection systems also require all readers to be constantly connected to a central server, which is not the case in some of the current and planned deployments of RFID tags where offline readers are used.

Tamper-proofing can be used to protect secret keys from attackers, but provides little help against hardware reverse-engineering because the structure of the circuits will always be preserved. The implementation, however, could be obfuscated to increase the complexity of the circuit detection. While we believe that obfuscations will not make our approach infeasible, we do not yet know to what degree obfuscations could increase the effort and cost required to reverse-engineer a circuit.

All low-cost cryptographic RFID tags are currently ill-suited for high security applications because they lack tamper-proofing and are vulnerable to relay attacks. In

these attacks, the communication between a legitimate reader and a valid card is relayed through a tunnel thereby giving the reader the false impression that the card is in its vicinity. No level of encryption can protect against relay attacks and new approaches such as distance bounding protocols are needed [8].

5 Conclusions

Reverse-engineering functionality from silicon implementations can be done cheaply, and can be automated to the point where even large chips are potential targets. This work demonstrates that the cost of finding the algorithm used in a hardware implementation is much lower than previously thought. Using template matching, algorithms can be recovered whose secrecy has so far provided a base for security claims. The security of embedded cryptography, therefore, must not rely on obscurity. Any algorithm given to users in form of hardware can be disclosed even when no software implementation exists and black-box analysis is infeasible. Once the details of a cryptographic cipher become public, its security must rely entirely on good cryptographic design and sufficiently long secret keys.

The cryptographic strength of any security system depends on its weakest link. Besides the cryptographic structure of the cipher, weaknesses can arise from protocol flaws, weak random numbers, or side channels. When random numbers are weak and the user identification is not properly mixed into the secret state, codebooks can be pre-computed that lead to attacks that are much more efficient than brute force. In the case of the Mifare Classic cards, the average attack cost shrinks from several hours to minutes. Their cryptographic protection is hence insufficient even for low-valued transactions.

The question remains open as to whether security can be achieved within the size of the Mifare Crypto-1 cipher. The area of less than 500 gates may be too small to even hold a sufficiently large state, regardless of the circuits needed for the complex operations required for strong ciphers.

Acknowledgments. This work was partially funded by the National Science Foundation through the CyberTrust program, award CNS 0627527. Any opinions, findings and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect those of the National Science Foundation.

References

[1] Ross Anderson. *A5*. Post to *sci.crypt*, 17 June 1994.

[2] L. R. Avery, J. S. Crabbe, S. Al Sofi, H. Ahmed, J. R. A. Cleaver, D. J. Weaver. Reverse Engineering Complex Application-Specific Integrated Circuits (ASICs). In *Diminishing Manufacturing Sources and Material Shortages Conference*, 2002.

[3] Andrey Bogdanov. Attacks on the KeeLoq Block Cipher and Authentication Systems. In *RFIDSec*, 2007.

[4] Stephen C. Bono, Matthew Green, Adam Stubblefield, Ari Juels, Aviel D. Rubin, and Michael Szydlo. Security Analysis of a Cryptographically-Enabled RFID Device. In *USENIX Security Symposium*, 2005.

[5] Harvey G. Cragon. *From Fish to Colossus: How the German Lorenz Cipher was Broken at Bletchley Park*. Cragon Books, 2003.

[6] Electronic Frontier Foundation. Cracking DES. In *Secrets of Encryption Research, Wiretap Politics & Chip Design*, O'Reilly & Associates Inc., 1998.

[7] Martin Feldhofer, Sandra Dominikus, and Johannes Wolkerstorfer. Strong Authentication for RFID Systems using the AES Algorithm. In *Workshop on Cryptographic Hardware and Embedded Systems*, 2004.

[8] Gerhard P. Hancke and Markus G. Kuhn. An RFID Distance Bounding Protocol. In *SecureComm*, 2005.

[9] Martin E. Hellman. A Cryptanalytic Time-Memory Trade-Off. In *IEEE Transactions on Information Theory*, 1980.

[10] Daniel E. Holcomb, Wayne P. Burleson, and Kevin Fu. Initial SRAM state as a Fingerprint and Source of True Random Numbers for RFID Tags. In *RFIDSec*, 2007.

[11] Seokhie Hong, Deukjo Hong, Youngdai Ko, Donghoon Chang, Wonil Lee, and Sangjin Lee. Differential Cryptanalysis of TEA and XTEA. In *International Conference on Information Security and Cryptology*, 2003.

[12] Pasin Israsena. Securing Ubiquitous and Low-cost RFID Using Tiny Encryption Algorithm. In *International Symposium on Wireless Pervasive Computing*, 2006.

[13] Auguste Kerckhoffs. La Cryptographie Militaire. In *Journal des Sciences Militaires*,

1883.

[14] J. P. Lewis. Fast Normalized Cross-Correlation. In *Vison Interface*, 1995.

[15] Philippe Oechslin. Making a Faster Cryptanalytic Time-Memory Trade-Off. In *Crypto*, 2003.

[16] NXP Semiconductors. *Philips Semiconductors leads contactless smart card market*, 2006.

[17] Andrew Tanenbaum. *News Summary of Broken Dutch Public Transit Card*. `www.cs.vu.nl/~ast/ov-chip-card`

[18] David J. Wheeler and Roger M. Needham. TEA, a Tiny Encryption Algorithm. In *Fast Software Encryption*, 1994.

Practical Attacks against the MSP430 BSL*

[Work in Progress]

Travis Goodspeed
1933 Black Oak Street
Jefferson City, TN, USA
travis@radiantmachines.com

ABSTRACT

This paper presents a side-channel timing attack against the MSP430 serial bootstrap loader (BSL), extending a theoretical attack with the details required for a practical implementation. Also investigated is the use of voltage glitching to attack a disabled BSL.

1. SUMMARY

The Texas Instruments MSP430 low-power microcontroller is used in many medical, industrial, and consumer devices. It may be programmed by JTAG or a serial bootstrap loader (BSL) which resides in masked ROM.

Recent versions of the BSL may be disabled by setting a value in flash memory. When enabled, the BSL is protected by a 32-byte password. If these access controls are circumvented, a device's firmware may be extracted or replaced.

In many versions of the MSP430, a password comparison routine suffers from unbalanced timing, such that processing an incorrect password takes two clock cycles longer than a correct byte. By observing external timing, it is possible to determine the correctness of individual bytes, drastically reducing the amount of time required to guess a password.[3]

This vulnerability had been previously demonstrated by the author in simulation, but it is in this paper that a practical implementation of the attack is first disclosed. Further, some early results in the use of voltage glitching attacks against the BSL are presented.

2. SERIAL BOOTSTRAP LOADER (BSL)

The BSL of the MSP430 resides in masked ROM. If an entry sequence is performed, as is depicted in Figure 1, the BSL–rather than the user application–is run. When two rising edges are observed on the TEST pin preceding the rising edge of the -RST pin that power on the chip, the BSL begins to execute instead of the user-defined application program. For those chips with dedicated JTAG pins, the same sequence is the same except that falling edges are sent on the TCK pin.[4]

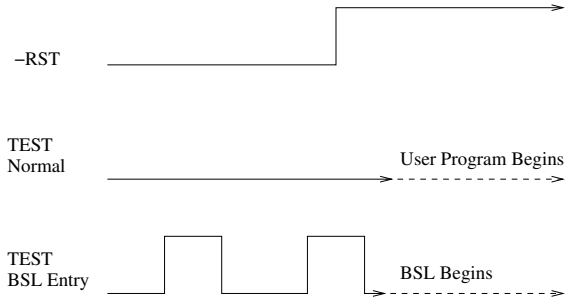

Figure 1: BSL Entry Sequence (Chips w/ Shared JTAG Pins)

As the BSL continues to function after the JTAG fuse has been blown, it is often used to allow for write-only updates without exposing internal memory to a casual attacker. For the same reason, it is a valuable attack vector. Each firmware image contains a password, and without that password little more is allowed by the BSL than erasing all of memory.

Once the BSL has loaded, commands are accepted through a bit-banged serial port. While there are many commands, the one of interest here is RX Password, which must preceed any attempt to read (TX Data) or write (RX Data) memory. Mass Erase, which bulk-erases all of memory, requires no password.

3. IVT AND PASSWORD

The BSL password is the Interrupt Vector Table (IVT) of the chip, which resides at the top of memory and is composed of sixteen 16-bit pointers to interrupt handlers. Of these 256 bits, the authors of [1] conclude that 40 are random. They then calculate that a brute force would take 128 years for a guaranteed break. This has since been reduced to 32 years in [2] by use of the Change Baud command. There might be room for further reduction, but the time required will never be so short as to be practical. Further, the method used to reduce the brute forcing time to the order of decades is only applicable to versions 1.60 and 1.61 of the BSL.

*Continuation of [3], presented by the author at Black Hat USA 2008

To appear in the Twenty-Fifth Chaos Communications Congress, December, 2008, Berlin, Germany.

```
d50:  jz 0xd56
d52:  bis #64,r11
d56:  dec r7
```

Figure 2: Byte Comparison in BSL 2.12

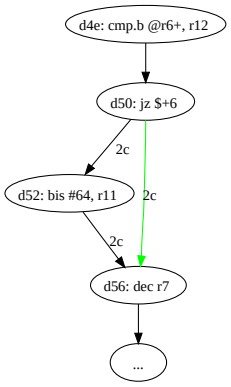

Figure 3: Control Flow of BSL 2.12 Comparison

Version 1.61 also carefully balances its timing to prevent the class of attacks presented in this paper. A cleaner solution would have been to bit-wise OR the XORing of each pair of bytes, as such a value being non-zero implies that the passwords do not match.

$$IVT = IVT' \iff \sum_{b \in IVT} b \oplus b' = 0$$

Version 2.12's comparison routine, as shown in Figure 2, suffers from unbalanced timing.[1] It is unbalanced in that one branch takes two cycles longer than the other to execute. As this code is part of a loop and the longer path is that of an incorrect byte, the timing of this program will be retarded by two cycles for every incorrect byte. The control-flow diagram in Figure 3 shows this graphically. Compare this to the invulnerable code of BSL 2.01 in Figure 4, which is has balanced timing.

[1]BSL 2.12 from the MSP430FG4618/G is used as an example throughout this paper. Others, such as 1.30 from the MSP430F1101A, are also vulnerable.

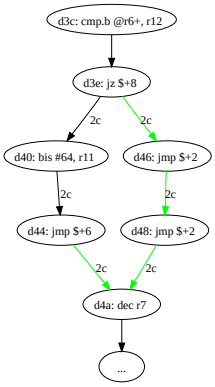

Figure 4: Control Flow of BSL 2.01 Comparison

4. SIMULATION

To demonstrate this in simulation, the author has written a C program for the MSP430 that wraps the BSL within an MSP430 simulator. The image was run 256 times, guessing passwords of every possible byte repeated. Timing was observed and recorded.

Shortened runtimes were found for repetitions of 0x00, 0x11, and 0x3A. Compared to an average (mode) runtime of 6543 cycles, a password of 0x00 repeated took only 6541 cycles to complete, a difference of 2 cycles. A password of 0x11 repeated took 6511 cycles, while 0x3A repeated took 6513 cycles. Thus the offsets were as shown in Table 1.

Guess	Cycles	Δ	$\Delta/2$
00*	6541	2	1
11*	6511	32	16
3A*	6513	30	15
all others	6543	0	0

Table 1: Runtimes of MSP430 BSL Wrapper

The rightmost column of the table gives the frequencies of each byte within the BSL. There must be a single 0x00, sixteen 0x11, and fifteen 0x3A bytes. As the less significant byte, being of an aligned instruction address, must be even, 0x11 is likely the more significant byte of each of 16 fields. Thus we have fifteen vectors of 0x113A and one of 0x1100. As the reset vector always points to the bottom of flash, it is 0x1100 and the rest are 0x113A. The BSL password for the image is shown in Table 2.

While this has demonstrated that the comparison routine itself has non-standard timing, more is required to break the BSL in practice. In particular, as the next section explains, any shift in timing will be hidden from measurement if the victim chip should wait for a start bit of a serial frame.

5. EXPLOITATION

The BSL runs at 1MHz until clocked higher[2], and a modern MSP430 can be clocked as high at 25MHz. Therefore, a 16MHz MSP430F2274 is quite capable of the timing necessary to break the password of a vulnerable chip. To that end, the author has designed a number of 'BSLCracker' boards around this chip which attack the BSL.

There are some complications, however. First, the BSL's timing is not hard-coded, but rather comes from a tare rou-

[2]By the Change Baud command, password protected after 1.61.

0x1100	0x113A
0x113A	0x113A
0x113A	0x113A
0x113A	0x113A
0x113A	0x113A
0x113A	0x113A
0x113A	0x113A
0x113A	0x113A

Table 2: Password of MSP430 BSL Wrapper

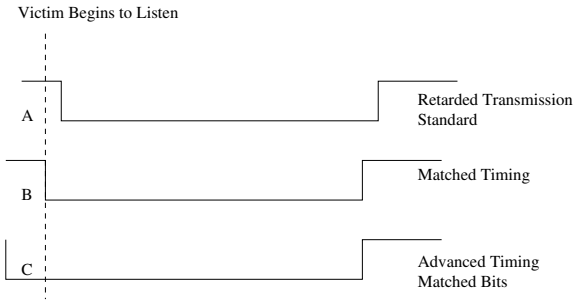

Figure 5: Three ways to say '0x80'.

tine.[3] This routine calibrates the bit-banging serial port handler by observing the timing of a header byte received by the MSP430. This header byte, 0x80, may be sent at an odd baud rate, something other than the standard 9600 baud that the BSL expects.

As Figure 5 demonstrates, there are three distinct ways in which a byte may be transmitted in a half-duplex serial port. (a) If the victim begins to listen before the start bit begins, as intended for reliable communication, all timing information is lost. There will then be no timing variance to observe. (b) If the victim begins to listen immediately after the beginning of the start bit, communication remains reliable and timing information is preserved. The attacker, however, would need to minimize the delay between the falling edge of the start bit and the victim's beginning to listen. (c) For this reason, it is practical to drop the start bit early, then gamble on the moment at which the victim begins to observe the byte. In this way, timing information is preserved whether the victim's timing is advanced or retarded from the estimate.

The technique of the preceding paragraph must be employed from the first byte of the password guess until the conclusion of the command sequence. Thus, as a checksum proceeds the password, checksum bytes must employ the same technique to avoid destroying timing information. The same is not true of header bytes, which precede the password and are thus irrelevant to the timing being measured.

6. LOCKOUTS, SELF DESTRUCTION

There are two self-protection features which have been added to the BSL in recent versions. The first is a lock-out, whereby the BSL can be disabled by placing 0xAA55 into the location named BSLKEY. The second is a self-destruct feature, where recent MSP430 chips will erase all of memory in the case of a first incorrect password attempt. The author is presently experimenting with a few methods of bypassing these restrictions.

A disabled BSL may be bypassed by voltage glitching, a technique borrowed from Smart Card 'Unlooper' technology. An R/C circuit is charged to a voltage which is significantly less than the minimum required by the victim chip. If this is timed properly, faults may be introduced into the behavior of that chip, such as the skipping of a register write-back. A scope recording of such a glitch is presented in Figure 6,

[3]See 0xE86 of BSL 2.12 from the MSP430FG4618/G.

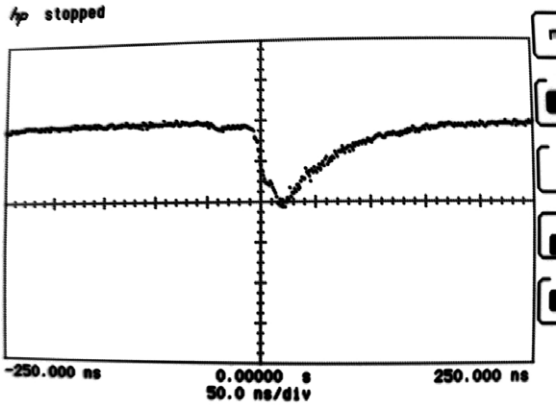

Figure 6: 45ns Voltage Glitch

```
c0c:  cmp #0aa55, &0xffbe
c12:  jz 0xc12
```

Figure 7: Disabled BSL Check, from BSL2.12

taken from the glitching of an MSP430 on the author's business card.

Reliable glitching is quite unnecessary to enter a disabled BSL, as can be seen by the code in Figure 7. When the BSL has been disabled by placing 0xAA55 into 0xFFBE, the CPU will repeatedly execute the jump instruction at 0xC12 until the watchdog timer resets the chip to the user application. During each execution of this instruction, a successful glitch will continue execution at 0xC14, rather than returning to re-execute the same instruction. An unsuccessful glitch might reset the chip, or it might have no effect. In the latter case, the attacker is free to try again on the next instruction. Should glitching become more reliable, it might be used to skip less often executed instructions within the BSL. For example, the password comparison routine might be exited after the first comparison, such that only the first byte need be correct.

While glitching register write-back is certainly the most impressive result of a changing supply voltage, there's a much less impressive effect that can be gained by setting the voltage above the minimum required for CPU operation yet beneath the minimum required for erasing flash memory. Attacks of this sort against flash memory have yet to be thoroughly investigated by the author.

7. HARDWARE AND SOFTWARE

A schematic diagram of BSLC30–the first attempt at the third major revision–is presented in Figure 9 and a photograph in Figure 8. Resistors and capacitors are added to test points for voltage glitching using the 74HC4053 MUX gates. These analog, bi-directional MUXes cut off all I/O traffic to the victim and drop voltage to a fraction of the minimum required for operation. Later revisions of the BSLC will replace the MSP430F2274 with another MSP430, one that contains a Digital to Analog converter more accurately selecting a target voltage. Further, some sort of level converter will be added to facilitate running the victim at an

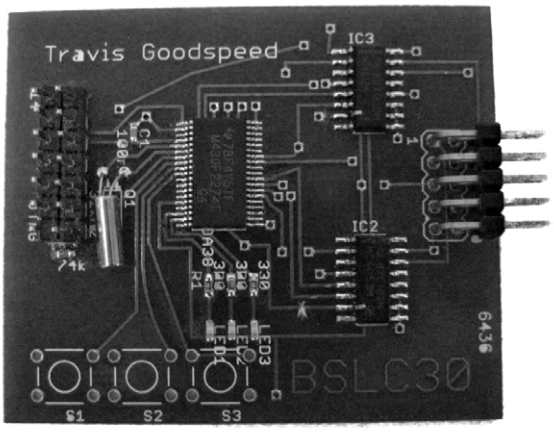

Figure 8: BSLCracker 3.0

arbitrary voltage during continued communication.

Neither the hardware nor the software of the author's implementation is complete, and further work is necessary to reliably attack the BSL. It is expected that as the software matures, the timing attack will give way to voltage glitching which–while more difficult to calibrate–is potentially effective at bypassing the access restrictions of every BSL version, even those with balanced timing.

8. RESULTS

By use of the BSLC and similar apparatus, the author has experimentally confirmed variable timing of the password-comparison routine in two of the vulnerable BSL versions, breaking a simple password by manual interpretation of timing results, similar to the method used in simulation. Further experiments have confirmed the vulnerability of the MSP430 to voltage glitching attacks by skipping over a "jz $+0" loop, both in isolation and to gain unauthorized entry to a disabled BSL.

Automated cracking of a BSL password has not yet been performed, but is expected within the next few months as the final firmware image is constructed. Glitching of MSP430 instructions other than tight loops will be the topic of forthcoming paper.

9. CONCLUSION

A side-channel timing attack against the MSP430 has been presented in sufficient detail for a practical implementation of attack. A method for using voltage glitching to enter a disabled BSL has also been presented, with citations of the code being glitched. Further, a brief outline of the author's implementation of an attack tool has been presented.

Versions of the BSL prior to 1.61 are forfeit to the timing attack, while versions after 2.01 are only vulnerable when the self-destruction feature has been disabled. Disabling the 2.x BSL by the BSLKEY flag is ineffective, and an MSP430 with a disabled BSL ought to still have a randomized password, such as by Alexander Becher's IVT randomization script.[2] As a disabled BSL erases flash memory on the first failed attempt at authentication, the timing vulnerability of recent MSP430 revisions is less serious than that of prior versions.

So far as the results of this research affect security, it should be emphasized that few if any general purpose microcontrollers are designed to defend themselves against a motivated attacker. The MSP430's insecurity, to both timing and voltage glitching attacks, does not in itself imply that similar chips from competing manufacturers are less vulnerable.

10. REFERENCES

[1] A. Becher, Z. Benenson, and M. Dornseif. Tampering with motes: Real-world physical attacks on wireless sensor networks. In *SPC 2006*, pages 104–118.

[2] T. Goodspeed. MSP430 BSL passwords: Brute force estimates and defenses, June 2008.

[3] T. Goodspeed. A side-channel timing attack of the MSP430 BSL. Black Hat USA, August 2008.

[4] S. Schauer. Features of the MSP430 bootstrap loader. TI Application Report SLAA089D, August 2006.

All code citations are from BSL 2.12 of the MSP430FG4618, Revision G, a product of Texas Instruments.

APPENDIX
A. BSL VERSIONS

Table 3 lists the BSL versions which have–and have not–been sampled by the author for vulnerability to the timing attack. All are presumed to be vulnerable to voltage glitching.

The BSL version is contained in memory at 0xFFA as a pair of bytes to be read visually. '0x02 0x12' therefore indicated version 2.12. The BSL itself spans from 0xC00 to 0x1000.

Version	MSP430	Timing
1.10		?
1.30	F1101A	Vulnerable
1.40		?
1.60		?
1.61	F1612	Invulnerable
2.01	F2274	Invulnerable
2.12	FG4618	Vulnerable
>2.12	?	?

Table 3: BSL Vulnerability by Version

B. BSL SYMBOL TABLE

It is expected that Table 4 will be valuable to those attempting to reverse engineer the MSP430 BSL from machine code. The values refer to the masked ROM of the MSP430FG4618, Revision G, from 0xC00 to 0x1000.

Address	Value/Type	Name
c00	0c06	Hard BSL Entry Point
c02	0c1e	Soft BSL Entry Point
c0c	code	Loop if Disabled
c1e	code	Watchdog Timer Disabled
c54	sub	Self-Erasure
ce0	sub	Set Main Offset
ce6	sub	Change Baud
cfe	sub	Load PC
d06	sub	Erase Segment
d0a	sub	Mass Erase
d2c	sub	Erase Check
d40	sub	RX Password
d50	sub	Password Byte Comparison
d80	sub	RX Data Block
dfa	sub	TX Data Block
e86	fn	Tare Bit Width
eec	fn	Write Byte
f54	fn	Read Byte
ff0	f46f	Chip ID
ffa	0212	Version in BCD

Table 4: BSL 2.12, FG4618/G

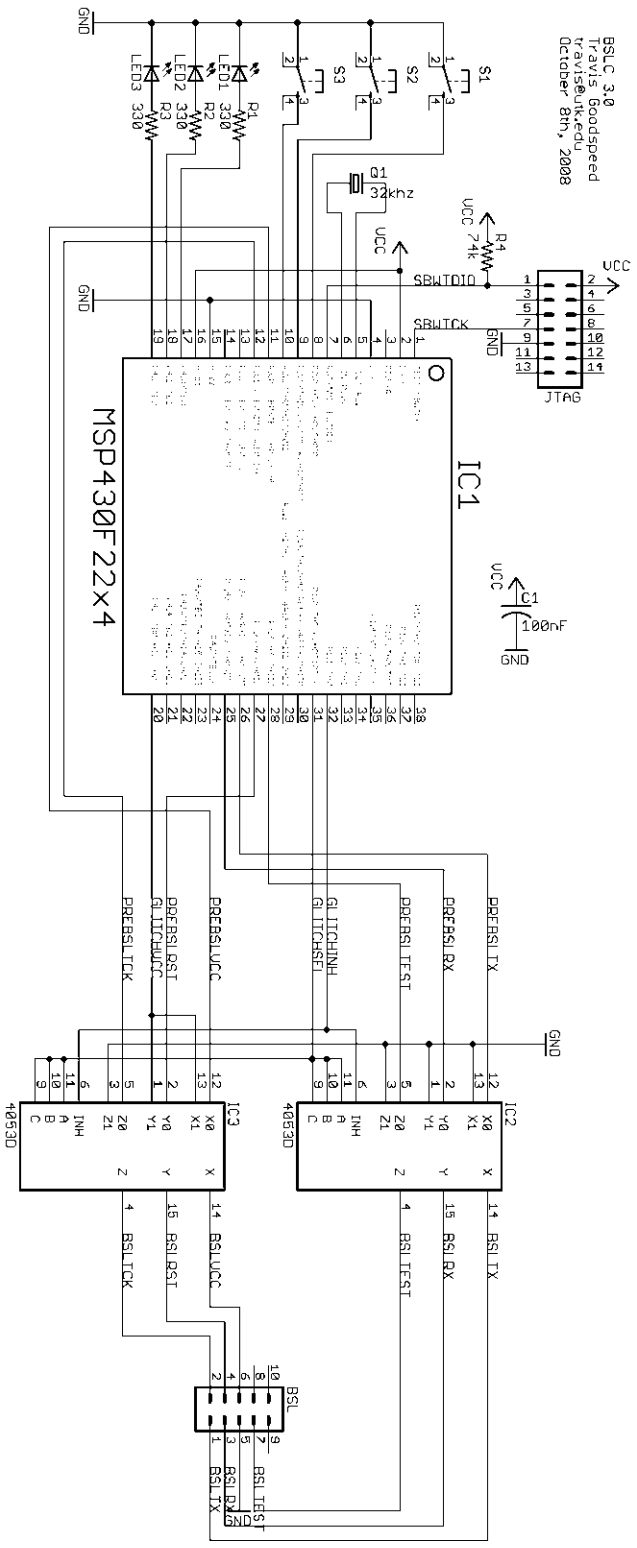

Figure 9: BSLCracker 3.0 Schematic

25C3: Full-Disk-Encryption Crash-Course
Everything to hide

Jürgen Pabel, CISSP
Akkaya Consulting GmbH

#0 Abstract

Crash course for software based Full-Disk-Encryption.

The technical architectures of full-disk-encryption solutions for Microsoft Windows and Linux are described. Technological differences between the two open-source projects TrueCrypt and DiskCryptor are explained. A wish-list with new features for these open-source Full-Disk-Encryption offerings wraps things up.

Readers should be familiar with general operating system concepts and have some software development experience.

#1 Introduction

Software based Full-Disk-Encryption („FDE“) solutions are usually employed to maintain the confidentiality of data stored on mobile computer systems. The established terminology is „data-at-rest protection“; it implies that data is protected against unauthorized access, even if an adversary possesses the storage media. Therefore, some of the most common risks to data confidentiality are addressed: loss and theft of mobile computer systems.

The encryption and decryption of data is completely transparent to computer users, which makes Full-Disk-Encryption solutions trivial to use. The only user interaction occurs in the Pre-Boot-Authentication („PBA“) environment, in which the computer user provides the cryptographic key in order to allow for the operating system to be started and data to be accessed.

#2 Architecture

I: Pre-Boot-Authentication

The Pre-Boot-Authentication environment is loaded before the operating system when the computer is powered-on. The primary purpose of the Pre-Boot-Authentication is to load the encryption key into memory and to initiate the boot process of the installed operating system. The following methods are some common examples of how the encryption key may be provided:

- by entering a password,
- by inserting a token (Smartcard, USB token, ...) or
- by loading a cryptographic key from a TPM chip.

Perhaps the most interesting technical aspect of the Pre-Boot-Authentication environment is the question about its storage location: if the entire disk is encrypted, how is the Pre-Boot-Authentication loaded? An obvious answer is to store the Pre-Boot-Authentication on an unencrypted media, like a bootable USB stick. However, booting from a third media imposes a physical dependency and it is therefore very common to store the Pre-Boot-Authentication in:

- a dedicated (unencrypted) boot partition,
- unallocated (not assigned to any partition) or „hidden“ sectors (Host-Protected-Area),
- sectors assigned to a partition but not bound to a filesystem (ie: filesystem was shrunk),
- an encrypted filesystem, but (only) sectors containing the PBA are left unencrypted or
- a designated boot area on the filesystem (if supported by the filesystem, like NTFS [0]).

The established approach on Linux is to use a dedicated boot partition (usually mounted as /boot), which is left unencrypted and loads the Linux kernel. An accompanying RAM-disk contains the Pre-Boot-Authentication environment, which reads in the encryption key and mounts the encrypted filesystem(s) before the boot process continues as usual (init process). Technically, this approach does not strictly meet the term „Pre-Boot-Authentication“ because the Linux kernel is already booted and used to obtain the cryptographic key.

Microsoft Windows requires a „real“ Pre-Boot-Authentication environment due to its boot architecture: a (usually rather small and product-specific) Pre-Boot-Authentication environment is loaded from any of the aforementioned storage locations instead of the NT boot loader (NTLDR). The Pre-Boot-Authentication environment reads in the encryption key, hooks interrupt 13h (INT13)[1] with a cryptographic function in order to allow transparent disk access via INT13, before passing control to NTLDR for loading the Microsoft Windows kernel (NTLDR uses INT13 for disk access). NTLDR loads the Microsoft Windows kernel, including all configured drivers, before the execution control is given to the kernel.

II: Encryption driver

Microsoft Windows disk drivers implement disk access without using INT13, therefore all encryption and decryption of data must occur within the kernel. The device driver architecture for Microsoft Windows includes a component called the „lower filter driver“ [2]. Encryption drivers for Microsoft Windows are usually registered as a lower filter driver for the disk drivers. The encryption driver initializes itself by copying the encryption key used by the INT13 hook function.

The mechanism of choice among Linux distributions for Full-Disk-Encryption is via the „device mapper“ [3] subsystem: a virtual device is mapped over the target device. The cryptographic algorithm is implemented in the virtual device driver and delegates all I/O operations to the underlying device driver. Loopback drivers and cryptographic filesystems are other alternatives for disk encryption on Linux. Yet another approach for implementing disk encryption on Linux would be to implement a kernel driver which hooks into the block device's function pointer (like `ide_driver_t.do_request` for IDE disks).

III: Initial disk encryption

Initial disk encryption for Linux mapped devices requires for the filesystem to be created anew, encrypting an existing filesystem is not possible (yet?). Solutions for Microsoft Windows are one step ahead in this respect, encrypting an existing filesystem („in-place encryption“) is commonplace – even among open-source implementations (refer to section #4/III).

The initial disk encryption is usually managed by a dedicated service process: it continuously loads, encrypts and saves disk sectors until a disk is fully encrypted. The initial encryption phase usually takes several hours to complete, depending on factors like CPU speed and disk size. Using the computer system during the initial encryption phase is possible but performance is often severely limited due to the constant use of the CPU and continuous disk I/O. The only performance control available in (almost) all Full-Disk-Encryption solutions governs CPU limits. However, todays CPU performance levels and multi-core architectures make it much more adequate to employ I/O limiting controls in order to more precisely define the acceptable performance impact during the initial disk encryption phase.

#3 Threats

I: Passwords

Whenever password based authentication schemes are employed for the Pre-Boot-Authentication environment, it's of utmost importance to choose a strong password in order to avoid brute-force attacks on the cryptographic key. Randomly chosen or generated passwords with upper- and lower-case letters and digits add roughly 6 bit of entropy to the key space per password character. It would therefore require a password with more than 40 characters in order to properly protect a 256 bit cryptographic key. While it may not be absolutely necessary to adopt such an extremely long password, it is recommended that passwords for encryption purposes should be significantly stronger than suggested by standard password recommendations.

II: Cold boot attack

A recently publicized security risk for computer systems protected by a Full-Disk-Encryption solution is called the „cold boot attack“ [4]. The attack targets powered-on (the operating system is loaded) computer systems by extracting the cryptographic key from RAM. It is therefore of utmost importance to power-off or hibernate computer systems if they are exposed to unauthorized physical access.

III: Coercion

Legal, physical and other forms of coercion are non-technical threats to data protected by cryptographic means. Plausible deniability is a technique which may offer protection from coercion in some circumstances by hiding an encrypted virtual disk (a „hidden volume“) inside another encrypted disk. Adversaries won't be

able to determine whether a hidden volume exists, unless external evidence – like a witness who knows about the existence of the hidden volume – indicates that a hidden volume must exist. TrueCrypt is the only solution that combines Full-Disk-Encryption protection with hidden volume capabilities [5].

#4 Products

I: History

Commercial Full-Disk-Encryption solutions have been around since at least the early 90's for Microsoft DOS. TrueCrypt comes to mind as the premier open-source offering, but one of the most desired features – the encryption of the system drive – has only been implemented in version 5, which was released in February 2008. The first open-source solution with support for system drive encryption on Microsoft Windows was released one month earlier: DiskCryptor [6].

II: Commercial Software

Many commercial Full-Disk-Encryption offerings exist. The most interesting/relevant solutions are:

Name	OS Support	Notes
CE-Infosys CompuSec	Windows & Linux	Free version (as in beer) available for personal and commercial use
CheckPoint FDE	Windows, OSX & Linux	Only FDE solution supporting system drive encryption on OSX
PGP WDE	Windows	Allows mounting of encrypted disk on another system („disk slaving“)
SafeNet ProtectDrive	Windows	Includes CD/DVD&USB encryption support
Secude FinallySecure	Windows	Linux based PBA
Utimaco SGE	Windows	PBA is obfuscated/obscured in order to harden analysis

Multi-user support for the Pre-Boot-Authentication environment is a common feature among commercial Full-Disk-Encryption solutions. Most solutions allow for the user's operating system login credentials to be synchronized automatically to the Pre-Boot-Authentication environment in order to not require users to remember yet another password. Automatic credential passing from the Pre-Boot-Authentication environment to the operating system login is also a commonly implemented feature: users are not required to authenticate at the Microsoft Windows login-screen, a software component automatically handles the login by using the credentials the user entered in the Pre-Boot-Environment.

III: Open-Source Software

TrueCrypt and DiskCryptor are open-source software projects and feature Full-Disk-Encryption support for Microsoft Windows. Neither have currently implemented any key or user management features and are therefore only of limited use in enterprise environments. TrueCrypt is a container based encryption solution at heart: it was primarily designed to use encrypted files as virtual disks. Support for encrypting real disks and partitions was added in version 5, but only for Microsoft Windows. Computer users with only modest computer experience will find TrueCrypt rather confusing due to its rather involved procedure for activating its Full-Disk-Encryption features. Technically versed computer users on the other hand value TrueCrypt's extensive cryptographic capabilities. DiskCryptor's user interface (which for the most part consists of the installer application) offers more of a traditional Full-Disk-Encryption experience from a computer user's perspective: the user is presented a dialog during installation to select the disk drives for encryption. The final installation step requires the user to set an encryption password, after which the selected disk drives are being initially encrypted.

TrueCrypt and DiskCryptor are entirely independent, but compatible, implementations: DiskCryptor adheres to the the TrueCrypt volume specifications. The TrueCrypt volume specifications requires a 512 byte volume header to be located before the encrypted volume. The header contains data required for mounting the encrypted volume (this applies to disks, partitions and file-back containers alike). File-backed containers

contain this header at offset 0, while the header is stored in the sector located just before the lower partition boundary for partition-back containers. However, this only works reliably for the first partition as the first partition always starts on (at least) sector 63, leaving sector 62 (and others) unused and available for the volume header. This workaround has a negative impact on TrueCrypt's Full-Disk-Encryption capabilities: only the first partition may be encrypted in-place as other partitions are unlikely to have a spare sector at their lower partition boundary. DiskCryptor on the other hand stores the volume header inside the first sector of the partition and therefore allows for any partition to be encrypted in place. However, this approach introduces two issues:

- it reduces the available partition size for the filesystem by one sector[1] and
- it requires relocating every sector in the partition.

DiskCryptor implements a custom filesystem shrinkage function for NTFS/FAT12/FAT16/FAT32 filesystems in order to shrink the filesystem on Microsoft Windows XP and 2003 as necessary. DiskCryptor running on newer versions of Microsoft Windows employs native filesystem shrinkage support, if available. The sector relocations are conducted during the initial encryption process, which reads data from an unencrypted sector and writes the encrypted data to the target sector. The sector shifting logic is implemented in DiskCryptor's lower-level filter driver and therefore hides all aspects about the sector shifting to upper layers. For example: if the filesystem driver initiates a read operation for sector `X` then DiskCryptor intercepts the request in the lower-level filter driver, loads sector `X+1` from disk and returns the decrypted data as sector `X`.

#5 Oddities

I: TPM Support

Current computer systems – especially laptops – are often equipped with TPM chips for secure cryptographic key storage and most commercial Full-Disk-Encryption products claim TPM support. However, none of the commercially available products support storing cryptographic keys on TPM chips. Instead, the TPM support is limited to TPM "binding": the ID of the TPM chip (its public key) is saved to the settings during installation and compared with the TPM chip's current ID. The comparison is executed either at the start of Microsoft Windows or for every user login; mismatches cause either an immediate shutdown or a rejected user login. This unintuitive use of TPM chips is easily explained: TPM chips are vendor-specific devices and require proprietary drivers. Full-Disk-Encryption vendors dread the cost of adapting proprietary drivers to their Pre-Boot-Authentication environments. TPM chips are therefore only used from within Microsoft Windows by calling the standardized TSS interface, which in turn calls the vendor drivers. BitLocker, the Full-Disk-Encryption solution provided with some versions of Microsoft Windows Vista, is the only Full-Disk-Encryption solution with support for storing cryptographic keys in TPM chips.

II: Multi-disk support

Support for encrypting multiple disks is a commonly named feature by Full-Disk-Encryption vendors. However, adding or removing disks once a Full-Disk-Encryption solution is installed is not so commonly supported. Several products allow the use of newly installed disks but do not support encrypting the new disks. The only way to encrypt additionally installed disks is through a work-around: decryption of all encrypted disks, software uninstallation, software re-installation and re-encryption of all drives from anew.

Various recovery and migration scenarios might make it necessary for an encrypted disk drive to be installed on another system as a data disk. This seemingly trivial procedure is also not well supported by the majority of Full-Disk-Encryption solutions. The most adequate concept for such tasks is called „disk slaving" (refer to section #3/II): it allows for an encrypted disk to be added to another computer system and for the encrypted disk to be used by loading the new disk's cryptographic key in the Pre-Boot-Authentication environment through an additional authentication step. A permanent migration of an encrypted disk to other system is therefore at least cumbersome, if not even impossible to manage for larger deployments.

1 A different header/footer is stored at the end of the partition, therefore the filesystem must be shrunk by about 30KB in total

#6 Open-Source wish-list

I: TrueCrypt compatible user and key management

The current TrueCrypt volume specifications are key agnostic. It is therefore impossible to implement any user and/or key management features. However, user and key management functions are a necessity for enterprise deployments. Breaking into the enterprise Full-Disk-Encryption segment with TrueCrypt compatible solutions therefore seems to require modifying the current TrueCrypt volume specifications. A specification compatible alternative would be to store all user and key management related data not within the volume header, but as part of the Pre-Boot-Authentication environment's configuration. The Pre-Boot-Authentication environment would than safeguard the encryption password for the encrypted volume, instead of protecting the volume's encryption key.

II: Alternative Pre-Boot-Authentication storage implementations

The Pre-Boot-Authentication storage locations listed in #2/I should be implemented for open-source projects in order to allow for flexible deployment solutions. External Pre-Boot-Authentication loading methods like USB and PXE boot implementations already exist and should be complimented by auxiliary key provisioning systems like loading the cryptographic key from a High-Security-Module (“HSM”).

#7 Conclusion

Full-Disk-Encryption is a fairly straight-forward technology, albeit different architectures have evolved for Microsoft Windows and Linux operating systems. TrueCrypt and DiskCryptor implemented Full-Disk-Encryption support for Microsoft Windows in 2008 and both are viable open-source solutions for personal computer systems. However, neither open-source solution currently offers critical management features, thus enterprise environments still turn to proprietary products for their Full-Disk-Encryption needs.

#A References

[0] http://technet.microsoft.com/en-us/library/cc781134.aspx

[1] http://en.wikipedia.org/wiki/INT_13

[2] http://msdn.microsoft.com/en-us/library/aa490241.aspx

[3] http://sources.redhat.com/dm/

[4] http://citp.princeton.edu/memory/

[5] http://www.truecrypt.org/hiddenvolume.php

[6] http://www.diskcryptor.de/

#B Author

Jürgen Pabel studied Computer Science and Information Assurance at Georgia Tech and Norwich University. He's one of the first German CISSPs and works as an IT-Security consultant at Akkaya Consulting GmbH in Cologne (Köln), Germany. His blog is located at http://blog.akkaya.de/jpabel/ and covers a broad range of IT topics with a slight focus on security.

#C Acknowledgments

Martin Modahl, Jens Neuhalfen, Pit Linnartz and another person (who wishes to remain anonymous) reviewed this paper: thank you guys!

OnionCat – A TOR-based Anonymous VPN

Bernhard R. Fischer

1024D/4DE0395F <*bf@abenteuerland.at*>

October 28, 2008

Abstract

TOR is an anonymizing network. It allows users to anonymously access Internet services. Its architecture guarantees that the real IP of users cannot be revealed in any way. TOR also provides so-called *Hidden Services*. Those are services which are hidden within the TOR network. This means that not only the user stays anonymous but also the service (destination). Hidden services have several benefits but unfortunately they are not very user-friendly and they have some protocol restrictions.

OnionCat manages to build a complete IP transparent VPN based on those hidden services, provides a simple well-known interface and has the potential to create an **anonymous global network** which could evolve to a feature- and information-rich network like we know the plain Internet today.

1 Introduction

TOR is an anonymizing network[1] within the Internet consisting of several nodes capable of forwarding TCP/IP sessions through it thereby hiding the origin at the destination end point. The location of a user i.e. his IP address is hidden at the remote site, e.g. the IP address of a user accessing a web service will not be revealed in the server's log files. Instead the IP of a random TOR exit node appears. An exit node is a TOR node at which a TCP/IP session leaves the TOR network.

This is a great feature because it improves a person's privacy especially if somebody resides under aggravating circumstances. Unfortunately TOR is not only used in the "right manner". Some one could also missuse it and if done right nobody will every discover who did wrong because even for the TOR network itself it's impossible to reveal the originating IP – deliberately it's a design feature. Always only the IP of the exit node appears in the public and depending on the law of the country where the exit node resides in it could lead to a law-enforced service shutdown or even something worse.

That's why people are usually not willing to run exit nodes.[2]

1.1 Hidden Services

The counterpart of a user who would like to hide his location is a service which should be hidden, i.e. a service which you know that it exists and you know how to access it but you don't know where it is. Basically it could be any type of service, e.g. a web service.

In plain old Internet this is more or less impossible because an IP address can always be traced back to an Internet provider and finally to a user or company. Hidden services [2] are services which exist only within the TOR network and of course they are also location hidden. That means they are not identified by an IP address but by an .onion-URL and the TOR network is able to find the right path to it but neither the user nor the TOR network can detect the IP address.[3]

[1] See the TOR project page www.torproject.org for general information.

[2] That's not the only reason but probably the most important one.

[3] Of course only if the service is configured correctly.

Beside location hiding there is a second great benefit: connections to hidden services do not leave the TOR network. No single exit node is needed and that's perfect because, as already mentioned, exit nodes are rare and because of that they are permanently traffic overloaded which results in a high latency.

Another benefit is that TOR guarantees end-to-end encryption from the client to the hidden service which is not true for connections to the Internet even when using TOR.[4]

That's why the use of hidden services is really interesting. Providing them increases the usability of TOR and the privacy of users and service providers.

1.2 The Problems

Unfortunately these .onion-URLs look like random numbers and characters – and in fact they are more or less random – which makes them really hard to remember, even harder than IP addresses because they have 16 digits.

But who really needs to remember IP addresses? Everybody uses names today. There is the domain name system (DNS) which resolves names to IP addresses. In plain Internet, name service is one of the most important ones. Nearly every user and every service uses names instead of IP addresses while using the network. The introduction of DNS – a distributed name resolution service – made the Internet more usable and opened it to a wider community.

But within TOR there is currently no resolving mechanism available for translation of names to .onion-URLs. Traditional DNS can not be used that easy because it is IP-based[5] (specifically the Internet class IN) and hidden services are .onion-URL based which can not be simply exchanged with IPs. From the TOR point of view those URLs are already names. Theoretically, an approach could be to use canonical names (CNAME) pointing to .onion-URLs but this would break authentication. Unlike IP addresses .onion-URLs provide authentication, i.e. using the .onion-URL a user can verify that a service really is the right hidden service and not any other one who pretends to be the right service. DNS basically does not interact with services that are associated with names, i.e. it cannot provide authentication as it is used for TOR and the security of users and services.

Even if someone deals with those .onion-URLs it's still not easy to use hidden services because the interface between an application and TOR is SOCKS [5] – a protocol for proxying TCP/IP. From a software modularity point-of-view it is a good idea to use SOCKS because it is a well standardized interface and many applications support it. But many do not! And every application that supports it needs user interaction to setup the right settings for SOCKS. A user should be able to **use hidden services without any differences to regular Internet services.**

Furthermore SOCKS version 4 only supports TCP/IP. There's no transport for UDP and other layer 4 protocols. Typically DNS is based on UDP which is an important protocol but cannot be used in combination with TOR.

2 Basic Considerations

Based on the previously mentioned considerations we suggest an **application interface on the IP layer**. With such an interface every protocol based on IP should be transportable.

On most operating systems such interfaces are available and provided by the kernel. On Linux, *BSD and other Unices there are kernel modules providing a layer 3 tunnel interface, usually called *TUN* device[6] but it's also available on Darwin (probably because of its BSD code base) and a similar model is available even on Windows.

OnionCat shall connect only to hidden services. As already mentioned they are addressed by .onion-URLs which are requested through SOCKS4a [4] and resolved by TOR itself. Obviously, because .onion is not a valid top level domain (tld) in Internet DNS.

Unfortunately, if using layer 3 which usually is IP, there's no such thing like a host name. We need a new IP-compatible addressing scheme for hidden services but this cannot be done by

[4] Unfortunately many users do not know this fact and believe that everything gets encrypted just because they use TOR.

[5] That's not a matter of design but a matter of the real (IP) world.

[6] It's similar to the TAP device which is a layer 2 interface.

just setting up a DNS service which resolves .onion names. It would break the authentication scheme of TOR's hidden services and it would imply user interaction again to configure a specific DNS that hosts the .onion tld.

TOR generates a .onion-URL [6] out of the public key of hidden services. It's exactly an 80 bit wide Base32 encoded string. Those 80 bits are one half of the SHA-1 hash of the public key. And that of course is derived by the private key. That's why those URLs are strongly related to the hidden service. We do not want to loose any of those bits because it would increase the probability for collision attacks thereby breaking the authentication scheme again[7] and it would deny reversability.

7 bits	1	40 bits	16 bits	64 bits
Prefix	L	Global ID	Subnet	Interface ID

Figure 1: Unique-local address format.

We use IPv6 addresses as a new addressing scheme for hidden services. IPv6 addresses are 128 bit wide, that's large enough for including 80 bits of an .onion-URL. According to RFC5156 [1] we use a network out of the *unique-local* address space. These are reserved for internal use in networks comparable to those of RFC1918 [7] of IPv4. As shown in Figure 1 the basic address format has a fixed minimum prefix length of at least 48 bits, additionally variable 16 bits for subnetting and 64 bits for the interface ID (host part). We don't need any subnet so we add the full subnet part to the interface part resulting in an 80 bits wide host part. The prefix length for those addresses is 48 bits.

48 bits	80 bits
FD87:D87E:EB43	.onion-URL

Figure 2: OnionCat addressing scheme.

According to RFC4193 [3] we set the "L"-bit to 1 and generated a global ID thus resulting in the new unique-local IPv6 prefix FD87:D87E:EB43::/48 – the OnionCat prefix. Address translation is easy by Base32-decoding the .onion-URL and inserting those 80 bits into the host part of the IPv6 address (see Figure 2). E.g. decoding 7fd22jhmqgfl45j6.onion leads to 0xf947ad24ec818abe753e. Putting this together with our OnionCat prefix it results in the IPv6 address fd87:d87e:eb43:f947:ad24:ec81:8abe:753e.

[7] With time it's getting even worse because there are known collisions in SHA-1 yet. [9]

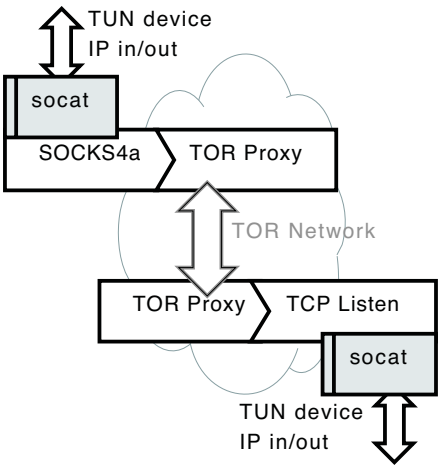

Figure 3: Socat hidden service connection.

Based on this addressing scheme we can now translate .onion-URLs to IPv6 addresses and vice versa.

OnionCat basically works similar to socat [8]. Socat[8] is a relay that handles all kinds of streams that are associate with a *file descriptor* (...in Unix everything is a file ;-). It has two ends each associated with a file descriptor and forwards data between them. For our application specifically interesting is the feature to forward data between a TUN device on one end and a SOCKS4a connection at the other end and at the opposite a TCP listener on one end and a TUN device again on the other end respectively (Figure 3). Now you could assign IP addresses to both ends and your IP-transparent point-to-point connection is ready. This setup has to be done manually.

Different to socat OnionCat **automatical connects through TOR** based on the .onion-URL related IPv6 addresses and it is able to build up **point-to-multipoint connections** because of its routing capability (Figure 4). The appropriate IPv6 address is assigned automatical to the TUN device which creates an entry to the kernel's IPv6 routing table. Hence, packets

[8] It's also father of the name "OnionCat".

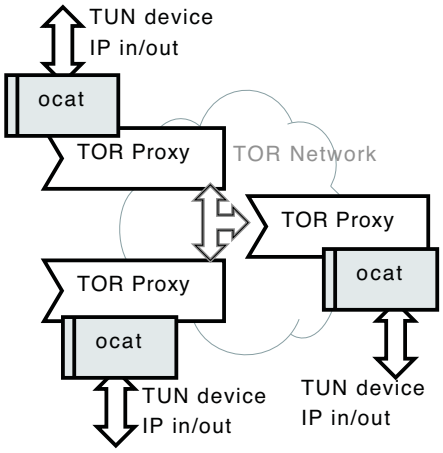

Figure 4: OnionCat connections.

are forwarded to OnionCat by the kernel without further interventions.

3 OnionCat Implementation

OnionCat is a multi-threaded application which basically receives packets on the TUN device and forwards them across the SOCKS4a connection through TOR and vice versa, once connections to remote OnionCat hidden services are established. Internally it maintains a *peer list*. Every *peer* is associated with an .onion-URL and its appropriate IPv6 address, the file descriptor of the TCP session (between OnionCat and TOR), an idle time value, some counters, and a defragmentation buffer.

Periodically the *Socket Cleaner* thread (see Figure 5) checks the idle times of the peers. If one exceeds the limit it is removed from the list and the Socket Receiver (see Section 3.2 below) thread is signalled that the peer list has changed.

The actions taken are different if packets are received either through the TUN device (outbound direction from a local point of view) or through a TCP session between OnionCat and the TOR proxy (inbound direction).

3.1 Outbound Direction

Packet reception on the TUN device is handled by the *TUN Receiver* thread (Figure 5). It extracts the destination IPv6 address of the incoming packet and *looks up* wether a peer with this address exists in the peer list or not. If so, it forwards the packet directly to the peer's file descriptor, updates his idle timer and continues receiving packets on the TUN device.

If there's no peer in the peer list it *initiates a new connection* by triggering a sleeping *SOCKS Connector* thread. The packet itself is queued for a while and gets forwarded by the *Packet Dequeuer* thread[9] after the peer is ready. Directly after the packet is queued the TUN Receiver continues receiving packets on the TUN device.

The SOCKS Connector spawns a new spare thread and tries to *connect* to the hidden service through the TOR proxy's SOCKS4a interface. It blocks until the connection is established by TOR. If it was successful it makes a new entry into the peer list, signals the Socket Receiver that the peer list has changed and terminates. If the connection failed it just terminates thereby dropping the request. The SOCKS Connector threads maintain a list of new peer requests.

3.2 Inbound Direction

Data reception from TOR is done by the *Socket Receiver* thread if connections are already established. If data is received it is appended to the defragmentation buffer of the appropriate peer. Every peer has its own defragmentation buffer. If the buffer contains at least one complete packet the source IPv6 address is extracted from the header and copied into this peer's address field if it is still empty (It will be explained shortly why this could occur). Then it forwards the packet to the TUN device and deletes it from the defragmentation buffer.

New incoming hidden service connections from TOR are handled by the *Socket Acceptor* thread. On program startup it creates a listening TCP socket and waits for connections (currently on default port 8060). Once a connection comes in it accepts it, creates a new entry in the peer list and continues accepting connections. The Socket Receiver is signalled that the peer list has changed.

At this time it does not know about the originating address (.onion-URL/IPv6) because those TCP sessions are always initiated by the local TOR proxy, hence, its source address is 127.0.0.1 (or ::1). Furthermore, it is just the transport, OnionCat (and every other hidden service) just

[9] This thread is *not* depicted in Figure 5.

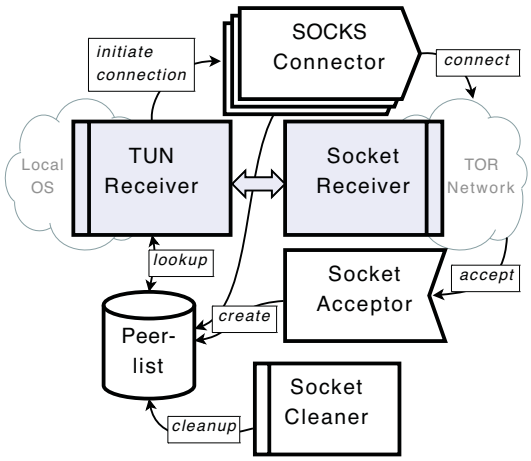

Figure 5: Internal structure.

uses the payload of those session. Outbound packets cannot be sent to this new peer as long as it is not identified. Identification happens immediately at reception of the first IPv6 packet (see above).[10]

3.3 OnionCat and IPv4

At a first glance it looks easy to do IP[11] forwarding if IPv6 does already work. But as mentioned in Section 2 hidden services are addressed by an 80 bits wide ID which we managed to convert to IPv6 addresses. Unfortunately that's not possible with IP addresses because they're only 32 bits wide. We could strip some bits off the .onion-URL and pack it into an IP address but sadly this type of conversion isn't reversible, but that's a requirement.

OnionCat does IP forwarding with a routing table which represents the glue between an IP address and an IPv6 address (and an .onion-URL respectively). Currently we use the least significant 24 bits of the IPv6 address and put them into the host part of the private network 10.0.0.0/8. The last 24 bits of the address fd87:d87e:eb43:f947:ad24:ec81:8abe:753e are 0xbe753e. Translating this to IP using our method results in 10.190.117.62. This IP address together with the netmask 255.0.0.0 is also assigned to the TUN device and an entry in the kernel routing table appears.[12] All packets with a destination therein will travel across the TUN device to the TUN Receiver thread (Figure 5).

The problem that now occurs is that the peer cannot be looked up in the peer list because the destination IP address cannot be reversed to an .onion-URL. As already mentioned, OnionCat maintains a second list which is a routing table with destination IP addresses, netmasks, and appropriate IPv6 gateways. On reception of an IP packet the TUN Receiver looks up an entry in the routing table and then further looks up the gateway address in the peer list and continuous as described in Section 3.1 or drops the packet if no routing entry exists. This routing table has to be setup manually.

4 Availability and Application

The source code of OnionCat can be downloaded at the current project home page www.abenteuerland.at/onioncat/. It is released under the GNU GPLv3 and is in *early* development state at the time of writing this paper. Currently it runs under Linux Kernels 2.6.x and 2.4.x, FreeBSD 6.x, OpenBSD 4.x, and Mac OS X 10.4 and 10.5, but maybe also under other operating systems. It's written portable as possible. The most ugly part is porting the TUN device initialization code (congratulations to the OpenVPN guys!).

Have a look at our project page for a description on OnionCat usage. We do not maintain a mailinglist yet, but we plan to do so. Announcements are currently done on the or-talk list.

The goal of what OnionCat is made for is to recreate the Internet on an anonymous basis: *AnoNet*. If everybody – this includes users and service providers, i.e. people, organizations, companies, etc., providing services – uses TOR and OnionCat, this could become reality.

For now it may be used for smaller user groups which need to exchange data basically with the same requirements as those using TOR but in a more transparent or more flexible way in respect to the underlying network protocols. Peo-

[10] This is a known security weakness but we'll find a solution.

[11] Subsequently I'll use the term IP instead of IPv4.

[12] Obviously, it may overlap other routing entries with subnets of 10.0.0.0/8 but OnionCat currently is in early development state and we don't care at the moment. We'll make this configurable in the future.

ple could setup private meeting rooms, chat relays, or similar services.

Of course, there are also dark sides. With OnionCat people can also do e.g. file sharing completely anonymous. But maybe this has a good side effect: if file sharing is done only within TOR, the exit nodes will become less overloaded.

5 Conclusion

In this paper we presented a method for making TOR's hidden services more user-friendly and transparent. This is done by insertion of OnionCat, a layer between client applications and the TOR proxy thereby lowering the access layer from TCP to IP. This change in layers also forces another addressing method for which we showed a deterministic reversible approach. By acting on the IP layer every protocol beside TCP can be transmitted without further circumstances across TOR.

OnionCat creates the major advantage of **using TOR's hidden services like usual IP hosts** on the Internet. Together with TOR, it has what it takes to build *AnoNet* – a perfect **anonymous Internet** within the Internet. This creates the interesting problem of anonymous not back-trackable payment methods.

During development and test phase we discovered some problems. OnionCat currently lacks authentication on incoming connections. It just uses the first incoming packet for identification.

IPv4 support is not very mature. OnionCat maintains an own routing table. It would be more comfortable if it shares the kernel routing table but we think that this might include portability issues.

The connection setup to a hidden service could last very long. Sometimes this takes one to two minutes. With this kind of transparency that OnionCat creates, also RTT measurements are easy, e.g. with the ping command in the simplest case. We observed RTTs between 1 and 30 seconds in the real TOR network[13] and unfortunately this seems not to be a matter of OnionCat.

Another problem seems to be that TOR is based on TCP. The circuits are built of concatenated TCP sessions. If creating TCP sessions through OnionCat it leads to "TCP-over-(TCP+TCP+...+TCP)". TCP uses algorithms to dynamically adapt to bandwith availability and this stack of dynamic systems could lead to very ugly behavior. Packet transmission sometimes looks like they are travelling through rubber bands which means that they are delivered in periodical occurring bulks.

But beside all that problems, we still believe in OnionCat and TOR and we think that this new kind of **anonymous VPN is a great benefit** for the people on this world.

[13] We maintain a private TOR network in our lab for test purposes.

References

[1] M. Blanchet. Special-Use IPv6 Addresses. RFC 5156 (Informational), April 2008.

[2] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation Onion Router. https://www.torproject.org/doc/

[3] R. Hinden and B. Haberman. Unique Local IPv6 Unicast Addresses. RFC 4193 (Proposed Standard), October 2005.

[4] Ying-Da Lee. SOCKS 4A: A Simple Extension to SOCKS 4 Protocol. http://ftp.icm.edu.pl/packages/socks/-socks4/SOCKS4A.protocol.

[5] Ying-Da Lee. SOCKS: A protocol for TCP proxy across firewalls. http://ftp.icm.edu.pl/packages/socks/-socks4/SOCKS4.protocol.

[6] The Tor Project. Tor Rendezvous Specification. http://www.torproject.org/svn/-trunk/doc/spec/rend-spec.txt, 2008.

[7] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. Address Allocation for Private Internets. RFC 1918 (Best Current Practice), February 1996.

[8] Gerhard Rieger. socat - Multipurpose relay. http://www.dest-unreach.org/socat/, 2007.

[9] Xiaoyun Wang, Yiqun Lisa Yin, and Hongbo Yu. Finding Collisions in the Full SHA-1. http://people.csail.mit.edu/yiqun/-SHA1AttackProceedingVersion.pdf, 2005.

Steel, Cast Iron and Concrete: Security Engineering for Real World Wireless Sensor Networks

Frank Stajano, Dan Cvrcek, and Matt Lewis

Computer Laboratory, University of Cambridge
15 JJ Thomson Av, CB3 0FD Cambridge, UK
frank.stajano@cl.cam.ac.uk, dc352@cl.cam.ac.uk,
ml400@cam.ac.uk

Abstract. What are the real security issues for a wireless sensor network (WSN) intended to monitor the structural health of a suspension bridge, a subway tunnel or a water distribution pipe? Could an attack on the sensor network cause the structure to collapse? How easy is it for civil engineers or other domain experts to build a secure WSN using commercially available hardware and software?
We answer these questions by conducting a qualitative risk assessment with bridge and subway tunnel operators and by conducting penetration testing on commonly available commercial WSN hardware and software, namely the Crossbow MICAz motes running TinyOS and XMesh and communicating over IEEE 802.15.4.

1 Introduction

We are interested in the practical security of real-world deployments of wireless sensor networks (WSN). The goal of our project is to develop reliable ways of monitoring the condition of large civil engineering structures such as bridges, subway tunnels, water pipes and sewers so that structural damage (due for example to corrosion, materials fatigue, overloading or shifting of surrounding soil) may be noticed and remedied before the structure weakens to the point of failure.

Traditionally, such monitoring is effected through periodic visual inspection. However the cost of accessing the structure for inspection is high, not just in terms of the effort required for the people in hard hats to reach the part of the structure to be monitored (think of the main steel cables of a suspension bridge or the middle section of a 3 km stretch of subway tunnel between two stations) but also in terms of the downtime of the structure for its regular users. This means that, unless there is reason to suspect a problem at a particular spot, routine inspections are infrequent (perhaps every few months) and cannot easily pick up new faults as they develop. This motivates our search for a solution based on a fixed network of sensors—wireless rather than wired in order to facilitate deployment in hard-to-access locations. Within the project, we (the three authors of this article) are responsible for the security engineering aspects of the wireless sensor network.

S.M. Bellovin et al. (Eds.): ACNS 2008, LNCS 5037, pp. 460–478, 2008.

From the security viewpoint, what is the problem to be solved? A preliminary challenge is to understand what damage could be caused by a failure (whether accidental or maliciously induced) of the sensor system. Would an attack on the sensor system just compromise the monitoring functionality, forcing the structure operators to fall back to manual inspection, or could it also have direct repercussions in the physical world, such as causing actual damage to the structure or to other nearby entities? If so, how? It is difficult to quantify such risks because there is essentially little or no prior experience of large civil engineering structures being monitored for many years by wireless sensors. Sometimes the structure operators themselves cannot imagine any serious security issues. It would be a mistake to launch into grand plans for encryption, key management and access control based just on what is fun for security researchers to do, rather than in response to recognised risks. So the consequences of security failures must be researched and understood, and the cost of security measures must be appropriate for the risks. The second challenge, assuming that the risk analysis requires and justifies such action, is to build a secure network of sensors using commercial off the shelf (COTS) hardware and software. Are current systems sold in a secure configuration? If not, is it still possible for non-experts to build a secure system out of such components? How? And how difficult is it?

In this paper we offer the following contributions. Firstly, we report on our initial qualitative risk assessment, carried out by interviewing the operating manager of a large suspension bridge and a contractor responsible for part of a large subway tunnel network (section 3). Secondly, and most significantly, we assess the practical security of the particular COTS system adopted by our team, the Crossbow MICAz motes running TinyOS or XMesh, together with the Stargate gateway: we design and implement a variety of attacks on this system and we report on the security problems we found, together with appropriate fixes where possible (section 4). As a further contribution to WSN security we ported the TinySec security library to the MICAz motes[1]. While some of our attacks exploit generally known vulnerabilities, others like selective jamming (section 4.3) and power exhaustion through routing table manipulation (section 4.5) are original and interesting in their own right. In section 4.3 we also demonstrate how an attacker can undetectably alter messages in an IEEE 802.15.4 radio environment. Finally, based on the experience we gained, we offer some architectural recommendations (section 5), independent of the particular hardware and software we used, that will help future teams design and deploy more secure WSN systems out of COTS hardware and software.

2 Scenario

2.1 Sensing

The purpose of a sensor network in our scenario is to monitor ambient conditions for hints that the structure may be deteriorating. Here is a non-exhaustive list of examples of what we monitor and why.

[1] The source code is available under GPL from
`http://www.winesinfrastructure.org/`

Bridges. The main cables of a 2 km suspension bridge are almost 1 m in diameter; each of them is actually a bundle of over ten thousand 5 mm steel wires, all anchored in strands into huge concrete blocks in underground chambers at either end of the bridge. We monitor temperature and humidity at various points in those chambers for signs of conditions that might lead to corrosion.

Some of the wires do break over the lifetime of the bridge owing to corrosion, defects and stress. A large safety factor is built in by design, to ensure that the main cables will still support the weight of the loaded bridge despite a number of inner wires having snapped; but monitoring the rate of breakages is very hard. A visual inspection, involving opening up the main cables with wedges, is extremely costly and disruptive; it will only spot breakages that are close to the inspection point and it may itself cause further breakages by stressing the cable. We will instead be using sensors based on acoustic monitoring of the "ping" sound made by a wire snap: with several synchronised sensors on the main cables, one may approximately locate the position of a breakage. Such sensors require a 100 kHz sampling rate, which imposes stringent performance constraints on the nodes.

Tunnels. A subway tunnel is essentially a hollow underground burrow whose walls are lined with large cast iron or concrete tiles. As the surrounding soil moves, settles and subsides, the cross-section of the tunnel deforms and the walls of the tunnel get damaged, perhaps developing cracks. We monitor existing cracks with sensors that measure the displacement across the edges of the crack. We also use inclinometer sensors to measure whether a given tile moves (by as little as a hundredth of a degree), indicating deformation of the tunnel that might lead to further cracks. We also monitor relative humidity, temperature and vibrations.

Water pipes. A large water distribution pipe can be almost 1 m in diameter. The opening and closing of valves causes pressure waves that stress the pipe and may eventually result in leaks. A leak in a large water pipe may discharge substantial amounts of water and cause a local flood in less than an hour. Prompt intervention is essential. We monitor water pressure at various points in the pipe and infer the presence and location of leaks with mathematical models.

2.2 Network Architecture

The standard architecture for this type of application is a three-tier wireless sensor network. At the bottom tier, the sensors are attached to nodes[2] that form a multi-hop ad-hoc network. Modern motes tend to use IEEE 802.15.4 as their communication standard[3] but some older motes may use unlicensed frequencies such as 868 or 916 MHz.

At the middle tier, the data measured by the motes is routed to a gateway, physically located near the nodes (e.g. inside the tunnel) because of obvious

[2] Also known as *motes* from the seminal "Smart Dust" paper [1].

[3] This standard is often colloquially indicated as ZigBee but, strictly speaking, most current motes do not implement the higher layers specified by ZigBee on top of the PHY and MAC defined by 802.15.4.

connectivity requirements. The gateway, usually an embedded Linux PC, may perform some preprocessing on the collected data. Communication between the gateway and the next tier is frequently implemented via GPRS, which is easy to deploy in isolated areas. However, GPRS is relatively slow and expensive and therefore may not be suitable for applications that require high data rates or always-on connectivity (as opposed to, say, overnight batch logging of a few data samples collected during the day). It also does not work in tunnels. Where ADSL is available, it is usually more convenient.

At the top tier, the data collected by the gateway(s) is aggregated into a database running on a central server (usually a PC) that will be accessed by various applications for visualisation and processing of the data.

3 Qualitative Risk Assessment

Since the value of our project to the security community comes primarily from its link to real-world installations, we started by interviewing two senior representatives of organisations that respectively operate a large suspension bridge and a system of underground tunnels. We wanted to elicit their perception of risks rather than relying on our own guesswork. This is an ongoing portion of the project and we next expect to interview a water distribution operator.

Direct consequences in the physical world. The disruptive potential of data modification or injection attacks is greatly amplified if the sensors are directly connected to actuators in a feedback loop (e.g. actuators close the valves as soon as a leak is reported). In such cases, altering bits in a radio message has a direct effect on the real world (a mains water pipe is closed down, stopping water delivery for a whole neighbourhood).

Our first round of questions to the bridge and subway operators therefore aimed to elicit whether the WSN would plausibly be linked to any actuators once the system reached maturity. For the bridge, the answer was a definite no: none of the measurements would have consequences requiring automatic and immediate action. Everything is reported to a control room from where human operators have direct visibility of the bridge. Any action, including such "soft" actions as activating road signs that lower the speed limit for motorists (something that is done in poor weather conditions such as high winds, frost and so on), is initiated manually by an operator who cross-checks the sensor readings with other clues (e.g. a windsock). For the subway tunnels the situation was more subtle, since operators cannot rely on a visual cross-check, but our contact could still not imagine a situation in which tunnel sensors would be directly attached to actuators without human intermediation. We are keen to interview a water operator, though, whose answer might be quite different in light of the leak scenario.

Confidentiality. We also asked whether the operators would be worried if sensor readings were not kept confidential. The bridge operator was not worried at all, since from his experience he did not expect any such reading ever to say anything

extraordinary or embarrassing. The tunnel operator was more sensitive and cautious: if a problem occurs he wants to hear about it first, and have a chance to address it before others (e.g. the press) know it occurred. A well designed WSN system should therefore be flexible; encryption should only be activated if confidentiality protection is worth more than its cost in energy consumption, key management and system complexity.

Short-Term Integrity. Concerning the threat of false negatives (attacker making us believe that there is no fault when there is one, thus stopping the structure owner from carrying out maintenance that might fix the problem, leading to aggravation of the problem with time and possibly partial or total collapse) we learnt from these interviews that, over the life of these structures up to this point (26 years for the bridge; 118 years for the underground tunnels), the number of serious structural incidents has been nil or very low. This implies that trying to take down the structure by hiding spontaneous faults that the sensor network would report is a strategy that might force the attacker to wait for a very long time! The situation might be different if the attacker also *caused* the damage, as well as was subsequently trying to hide it from the sensors, but if this were perceived as a serious threat then safeguards against causing physical damage would be much more of a priority than those against network tampering.

With the injection of false positives, instead, the sensors report damage (e.g. several wire snaps one after the other) that has not actually occurred. This forces the maintenance operators to waste time trying to locate and repair a non-existent fault, e.g. by opening up the main cable with wedges without finding anything. It may also force temporary closure of the structure while the problem is investigated. Therefore, false positives would be more disruptive, since unlike false negatives they could be triggered at will by the attacker. In any case the feedback was that any major alerts from the new system (WSN) would be viewed with some suspicion by the operators until the system had proved its worth. This was particularly true for the bridge operator, for whom the new system was a nice extra but not a necessity, whereas the subway operator had no "eyes" in the tunnel and was therefore more eager to accept and trust any technological development giving him greater monitoring power.

Medium-Term Integrity. The subway tunnel operators see value in analysing sensor data over the medium term (e.g. months) because they currently have no systems allowing remote monitoring of the state of the tunnels; and the tunnels themselves can only be physically inspected in the middle of the night when trains do not run. Systematic collection of data about the state of the tunnels would be very useful in budgeting for maintenance costs as it would allow more precise estimates of necessary works. Maintenance budget negotiations are currently based on very vague estimates derived from manual inspection of a small sample of accessible sections of the tunnels. Continuous monitoring would provide more reliable quantitative estimates that all parties involved would have an easier time accepting.

Long-Term Integrity. For structure operators, a significant perceived benefit of our automated sensing is the provision of decades-long monitoring logs that will

allow them to carry out previously impossible research on long term behaviour (and, inevitably, decay) of materials when plotted against influencing factors such as humidity, acidity, pollutants and stresses. The monitoring thus becomes particularly valuable when it allows reliable data collection over many years. One requirement is therefore that data be collected and stored in a sensibly designed non-proprietary format that can still be read after decades, under the assumption that any component of the physical implementation of the system will be replaced by something newer in due course. Another consequent requirement is on data integrity: corruption of sensor readings, especially if on a small enough scale as to go undetected at acquisition time, would invalidate the whole historical database and make the exercise useless. This requires an architecture in which, regardless of confidentiality, integrity of sensor readings is preserved at all costs.

Availability. We explained that it would be technically impossible to eliminate denial of service attacks on a wireless network and tried to understand their practical consequences for the operators. Apart from the financial loss of having wasted money on an unusable sensor system, even a complete jamming of the wireless network did not appear as a grave loss to the bridge operator, who did not expect ever to depend entirely on the WSN output; but it sounded somewhat more embarrassing for the subway operator, who anticipated his newly gained real-time "eyes" on potential cracks as a facility he would not want to do without.

Conclusions. The general conclusion from the interviews with the structure operators is that data integrity is the most important security property for this type of application. No direct link from sensors to actuators is envisaged in the two systems we discussed, so the main effects of an attack on the WSN are invalidation of collected data or incapacitation of the WSN system rather than damage to real world facilities. This is therefore not a very high risk setting. Having said that, there is still scope for attackers causing disruption to users of the structure insofar as denial of service or injection of false positives may lead the operators to close the bridge or the tunnel for safety reasons while the problem is investigated.

4 Attacks on a Real System

Inspired by the above-mentioned user concerns, we examined the available WSN technologies for vulnerabilities threatening the desired security properties. As we did that, we also found problems that the operators had not anticipated.

Our goal was to assess practical security of wireless sensor networks on a real, physical system, as opposed to just in theory or through simulations. So we targeted our attacks on the particular platform that our project adopted, namely the MICAz mote from Crossbow, running TinyOS v1.1 and XMesh from MoteWorks 2.0.F, together with the Stargate rev. 1.2 as a gateway.

A superficial reader might comment that, since we chose the components and assembled the system ourselves, any security holes we find only reflect on our own incompetence. On the contrary, the spirit of our investigation was to

imagine that a team of application experts (in this case civil engineers), assumed to be security-conscious but not security experts, puts together a system using COTS components, following the manufacturer's instructions and activating any recommended security features. We set out to assess the practical security of the resulting system and to suggest ways of improving it where appropriate.

Our limited budget and manpower would never have allowed us to carry out a comparative study of all commercially available WSN platforms to determine the most secure one, so that was never a goal. Nonetheless, we believe our results will interesting for users of other platforms too.

Each of the attacks or exploits described in this section has been carried out and validated on actual hardware. We report sufficient details to convince the reader that a vulnerability exists and has been exploited by us, but stop short of supplying malicious readers with a cookbook. We also describe how to fix the problem wherever possible. As a courtesy we supplied a copy of a preliminary version of this paper to Crossbow in September 2007, to give them a chance to release security patches based on our advisories.

4.1 Our Platform

The TinyOS operating system runs on various hardware platforms including MICA2, MICAz, Iris (Crossbow Inc.), Tmote (Moteiv), Intel Mote, Intel Mote 2 (Intel). It is a modular system that allows easy extension with drivers for new sensor boards or functionality.

TinyOS v1.1 appears to be the most commonly used version in practice: v2 exists but is not stable yet. TinyOS may be deployed as is, or in conjunction with a commercial derivative such as XMesh from Crossbow or Boomerang from Moteiv. Being an open source project, TinyOS is reasonably easy to analyse for stability and security issues. TinyOS v1.1 has been stable since September 2003, so it can be considered a fairly mature product.

We began by analysing TinyOS, focusing our attention primarily on cryptography and routing protocols. The first big surprise was that, while TinyOS ships with the cryptographic module TinySec [2], the latter can only be compiled for MICA2 motes and there is no implementation available for the current generation of motes with 802.15.4-compliant radio chips. This means that all the networks based on modern Crossbow, Moteiv or iMote devices are vulnerable to a slew of attacks from even relatively unskilled attackers. To address this deficit, we ported TinySec[4] to make it run with the latest 802.15.4 chips—namely the Texas Instruments CC2420 chip, used in most of the motes mentioned above—so that we could test our attacks on cryptographically secured networks. In the process of performing this port, we noticed that the TinySec MAC generation code uses a fixed block size of 7 bytes, while the underlying block cipher has a fixed block size of 8 bytes. While we do not believe that this causes a security exposure, our port corrects this behaviour to use the same block size

[4] See footnote 1.

as the underlying cipher[5]. Note that, as the TinySec authors themselves point out, TinySec provides no protection against replay attacks. Unfortunately this protection does not exist at any layer of the TinyOS radio stack either. In addition, since by design TinySec only supports a single key shared across the entire network, there is no protection against physical capture of one mote.

A general virtue (security-wise) of the TinyOS code is that it is very lean and spartan. The underlying system supplies minimal functionality, preferring to export simple primitives to applications. One of the results of this is that the basic TinyOS system seems to be fairly secure without making much apparent effort to be so; several times during the course of our analysis we tried out attacks that we thought might work, but were blocked by TinyOS' minimalist philosophy. In contrast to this, the commercial system we analysed, Crossbow XMesh, exports a very rich set of features to applications running on it and consequently exposes a much wider attack surface.

Concerning power consumption, a mote using its radio continuously would exhaust a pair of alkaline AA batteries in a couple of days (a couple of weeks for the expensive D lithium batteries we use). In normal use, with a duty cycle of around 1%, a mote lasts for several months.

4.2 Classes of Attacks

We chose not to concentrate on **physical attacks** [3] on the sensors and on the nodes attached to them, not because we think they are impossible[6] but because an attacker with physical access to the sensors could with comparable effort stage much more destructive attacks on the structure itself, for example by using explosives. We therefore focus on attacks on the communication systems, primarily the ad-hoc radio used by the sensor nodes but also the back-end link from gateway to central server.

We studied three broad types of attacks: **data payload attacks** that change the content of data packets; **network attacks** that affect the functionality of the network, for example by preventing communication, taking down specific links, modifying the routing topology or rewriting the firmware of a node; and **system attacks**, potentially the most damaging, in which the attacker exploits a vulnerability in one part of the system architecture (e.g. the wireless network) to gain control of other parts (e.g. the gateway or the central computer).

Attack mechanisms we employed included jamming (at various degrees of selectivity and at different layers in the stack), replay attacks, packet injection or corruption (where the injected or malformed packets were specifically crafted to probe for vulnerabilities or to trigger known vulnerabilities) and ACK spoofing.

[5] This divergence from a correct implementation will become a compatibility issue if end-to-end keys are used, as the MAC algorithm will have to be implemented on both the motes and the gateway. The NesC code for the motes cannot be compiled for the gateway, so its use of non-standard parameters hinders interoperability.

[6] On the contrary, with so many unsupervised nodes over a large area, we believe one cannot exclude that a few nodes may at some point be physically captured, even though attackers are unlikely ever to be able to take over a *majority* of the nodes.

4.3 Jamming

While it is well known that no one can prevent physical jamming of the radio channel, appropriate design decisions can reduce a system's vulnerability to denial of service (DoS) attacks and increase the cost of such attacks. In this section, we investigate the baseline case of an attacker using an unmodified COTS mote as the transmitter. It is understood that a more resourceful attacker, with a higher power transmitter, could cause greater damage.

Description. We used a MICAz mote to jam the communication between other motes by transmitting at the same time as them, thereby causing a collision. Unlike most previous jamming attacks, ours is selective and jams only a subset of the packets (e.g. those coming from a specified victim) while leaving others unaffected. The attack can therefore be used to selectively disconnect individual motes or whole regions from the network. It is also hard to detect because unaffected nodes do not notice anything unusual.

The algorithm is as follows:

1. We select criteria for messages to be jammed, e.g. senderID = 3 and messageType = `AM_MULTIHOP`.
2. We compile the code with the selected criteria into the attacking mote and deploy the mote in the vicinity of the victim.
3. When the attacking mote detects a transmission, it listens to just enough of the message to determine whether the packet meets its jamming criteria.
4. If the message meets the criteria, go to step 5. In our example, we have a match if byte 13 = 3 and byte 9 = `AM_MULTIHOP`. Loop back to step 3 if the criteria were not met.
5. The attacking mote switches the radio chip to transmit mode and jams the rest of the message by transmitting on the same channel.

Steps 3 and 5 are difficult for technical reasons. To start with, the CC2420 is a packet based radio, so in normal operating conditions the microcontroller is only informed of radio activity after a complete packet has arrived. This obviously makes step 3 difficult since, by the time we can tell that we should jam a packet, it has already been transmitted. We were able to overcome this difficulty by putting the radio chip into a debugging mode, where each bit received from the radio was sent to one of the microcontroller's input pins as it arrived.

Our second problem in implementing the attack arose from the tight timing requirements imposed by the data rate of the radio—250 kbps. The ATMega128 microcontroller is clocked at 7.37 MHz, effectively giving us 30 instructions to process each bit. This would not be too much of a problem were it not for the fact that invoking a hardware interrupt appears to take around 30 clock cycles, meaning that our interrupt handler (which gets invoked once per bit) was always too slow. However, having the first invocation of the interrupt handler busy-wait for the remaining bits in the frame allowed us to meet the timing requirements.

Finally, switching from receive mode to transmit mode takes some time, so if the frame was particularly short we would frequently miss the end of the packet by the time we started jamming. By iteratively optimising the code, we were able to reliably jam frames with as few as 5 bytes of data.

Empirical Results. The 802.15.4 standard implemented by the motes uses direct-sequence spread spectrum (DSSS) to increase resilience to noise and jamming. However, since 802.15.4 specifies the spreading code to use, our motes shared the spreading code with the target, thereby nullifying the protection offered by DSSS. We carried out several experiments to measure the success rate of jamming. The experiments used three motes: transmitter, receiver, and jammer. The transmitter sent out a packet every 200 ms and the receiver, connected to a laptop, forwarded the received messages to a laptop where they were logged. All three motes were positioned 45 cm above the ground.

The jamming appeared successful in an open space, but with some anomalies. All messages were jammed when the attacking mote was between the sender and the receiver. The initial results (see Figure 1) suggested that the angle defined by transmitter-receiver-jammer had a significant impact on the success of the jamming, but later experiments showed a lack of repeatability, with variations between 30 and 88% for a given position. Complete jamming (100% of frames jammed) was achieved in several configurations of the three motes.

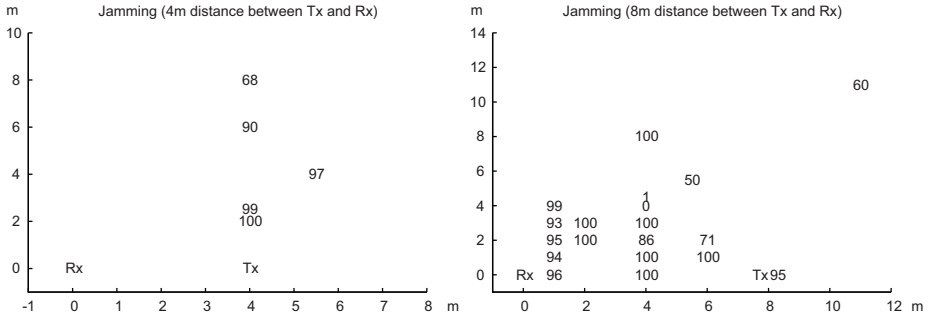

Fig. 1. Success rate of jamming depending on the position of the attacking mote. The Tx and Rx labels are the transmitting and receiving motes. The numbers 0–100 in the graphs denote the percentage of packets that were jammed when the attacking mote was in that 2D position relative to Tx and Rx (section 4.3).

We tried to find out why our results were so unpredictable, especially when the motes were close to the ground. A radio engineer told us that this is due to interference between the direct ray and the ray reflected from the ground, as the path difference between the two rays can be close to a half wavelength when the motes are on the ground. This results in destructive interference between the two rays, causing a weak signal. We think it may be some kind of voodoo.

Risks. The ability to selectively jam frames allows an attacker to:

- Make specific sensors deaf and/or mute;
- Prevent any sensor data from reaching the database, by jamming close to the gateway;
- Perform a man-in-the-middle attack by having a second mote listen for jammed frames and resend chosen ones with altered data.

The first two of these are simple DoS attacks and so may be written off by some as unimportant. The ability to alter the contents of any message while remaining undetected is rather more serious, as it gives the attacker complete control over the output of the network and therefore indirect control over all the systems controlled by the data reported by the sensor network.

Since the jammed frames are silently dropped by the radio chip, TinyOS is never notified. As such, a mote is unable to determine that a jamming attack is taking place unless its radio driver has been specifically modified to run in debug mode in order to detect this.

Underlying Cause. The vulnerability of a network to jamming as described here is simply a result of communicating over a shared medium to which adversaries have access (in this case radio).

Fix. To make it more difficult for an adversary to perform this type of selective jamming attack, the frame header could be moved to the end of the frame. This would mean that by the time an attacking node had decided to jam a frame, it would be too late. Unfortunately this defence is not perfect and succumbs to an attacker with two motes: one to jam every frame and one to resend any frames that do not meet the jamming criteria[7]. It is also expensive in terms of battery life, as each node must listen to and buffer the whole packet, instead of just the header, before knowing whether it was destined for itself. Detection of this type of jamming could be achieved by having one or more "observer" motes running in debug mode and listening for suspicious patterns of corrupted frames.

4.4 Counter Overflow Attack

Description. Another attack that can be mounted against a wireless network is a replay attack. This type of attack works even against an authenticated and encrypted network. One of the simplest protections against it is to use monotonically increasing counters to ensure the freshness of messages. TinyOS and TinySec use a 16 bit counter to distinguish messages, with TinySec using another counter as part of the encryption initialisation vector. Although TinySec explicitly does not offer replay protection, even if it did, the attack presented in this subsection would still work, unless TinySec increased the length of the counter. The network stack decides whether to accept a received message as valid according to the value of the counter in the message[8]. If the counter in the message is higher than the stored counter, the message is accepted. If it is lower but the difference is within an allowed range, the message is also accepted. Any

[7] We are not discussing all the details related to the necessity for the attacker to learn the whole content of every packet it jams in order to retransmit it if needed. This can be achieved in several ways, including jamming a part of the message that the attacker can reconstruct (such as the CRC, provided that there were no real errors), or jamming the message twice in different places, counting on the fact that the sending node will retransmit—although this latter strategy would be less stealthy.

[8] This is also used to compute the number of missed messages and subsequently the quality of the link with that neighbour.

other value causes the message to be rejected and the stored counter value to be zeroed. Each mote tracks the counter values of all its neighbouring motes.

For a replay attack to work, the receiving node needs to be ready to accept a message with the same counter value as the one we eavesdropped and intend to replay. We can achieve this by having the recipient's counter overflow and wrap around to the desired value. We thought of at least three methods to cause a counter overflow. First, we can inject fake messages that the target mote will forward, causing the target's counter to increment once per message we send. Second, we can use the routing table attacks described in section 4.5 to create a loop in the network: this causes a packet storm which eats up counter values. Finally, we can simply jam message acknowledgements from the target's parent as this will cause the target to retransmit each message several times.

Note that, if XMesh receives a message whose counter is set to zero, it bypasses the routing mechanisms used to verify the freshness of the message, obviating the need to overflow the counter. In addition, the receiving mote will not increment its stored counter, which means that the message has no effect on the link quality calculations. An attacker might exploit such "features".

Risks. When the message counter overflows we may initiate a replay attack with previously eavesdropped messages, even if these messages feature cryptographic integrity protection. This allows us to inject false negatives, false positives, or just slightly corrupted data to invalidate the long term monitoring. It also allows us to manipulate routing messages, even if they are authenticated. The message counters form the basis of many cryptographic mechanisms that rely on freshness, particularly in the absence of timestamps. This vulnerability would affect security mechanisms built from these cryptographic primitives.

Underlying Cause. Our three counter overflow methods rely respectively on the lack of source authentication, on the possibility of routing table manipulation (section 4.5), and on jamming (section 4.3). As such, the underlying causes of the counter overflow attack are the union of the underlying causes for these sub-attacks.

4.5 Routing Table Manipulation

Description. Nodes in an XMesh network populate their routing tables by listening to periodic broadcasts issued by their neighbours, which contain:

- The node's current parent;
- The path cost of sending a message from the node to the base station;
- A list of some of the node's neighbours and an estimate of how well it can receive messages from each.

When a node receives such a message from a neighbour, it updates its internal table of neighbours, which it consults every few minutes[9] to choose a new parent. All future messages are routed through this parent. The new parent appears to

[9] Every 30 seconds if the node is disconnected.

be chosen to minimise the path cost of sending messages back to the base station, with the link quality to each potential parent being a factor in the cost of the hop to the parent. The routing algorithm may thus be thought of as a distributed single-source shortest path algorithm.

Since all communication in the XMesh network is unauthenticated and unencrypted, an attacker can interfere with routing messages and thereby change the topology of the network to suit her whims. We created a Python module called `wsn` to read, modify, create and inject messages into the sensor network and we used it to implement and validate these routing attacks.

Risks. The ability to choose a topology for the network gives the attacker a lot of power. For example, she can cause all the traffic in the network to be routed through a mote she controls, allowing her to selectively drop or modify any message. Alternatively, she can cause a direct attack against the motes themselves by creating routing loops where, e.g. one route is $2 \mapsto 3 \mapsto 2 \ldots$ In such a configuration, nodes 2 and 3 will send messages back and forth, rapidly consuming their batteries. Using this method, we have been able to increase the peak rate at which nodes send messages to over 300/s with a mean value of around 25 messages per second. This can be used for a very powerful "sleep deprivation torture" [4] attack.

There is a low power implementation of XMesh that puts radio and microcontroller to sleep and they wake up only when necessary. However, any mote forwarding for at least one of its neighbours must wake up the radio chip regularly. The XMesh low power mode wakes up the radio chip roughly every 125 ms and it listens for a period of 1 ms. This means that the duty cycle is around 0.8%. The attack increases this value substantially to tens of percents.

Karlof and Wagner [5] discuss other routing attacks.

Underlying Cause. The problem here seems to be that the topology of the network is decided by the motes themselves while they do not have enough trusted information about the physical structure of the whole network to make sensible decisions. The ad hoc nature of the network appears to cause significant security issues, as everything received is trusted.

Fix. A natural solution to such trust issues would be to cryptographically authenticate messages sent by legitimate nodes. We could, for example, have all the nodes in the network share a key and use TinySec to authenticate and encrypt all messages sent over the network. This would mean that an attacker would not be able to forge routing messages without knowledge of the key. One problem with this scheme is that physical capture of a single node would reveal the network key and render the entire network susceptible to attack. Another problem is that authentic routing messages from one part of the network may still be recorded and replayed in other parts of the network, causing routing anomalies. We could stop this type of attack by using per-link keys (unique keys shared by each pair of neighbours), at the cost of complicating the key management process.

No cryptographic fix, however, would stop what has been called a "flagpole" attack, in which the attacker moves a victim mote up and down a flagpole in order to make other motes waste their battery updating their routing tables.

4.6 Over-the-air Programming

Description. One of the optional features that can be compiled into the motes making up an XMesh network is over-the-air programming (OTAP). This is a mechanism by which an entire network of motes can be reprogrammed remotely by sending them the code to execute. If OTAP is enabled on a network, no authentication is required to request a reprogram; an attacker able to send traffic to the network can therefore cause every mote to execute code of her choice.

Risk. The ability to reprogram motes allows an attacker to entirely control all data sent by the network. This means that she would be able to choose whether to send data from the sensors, what the apparent readings of the sensors were and how often the readings were reported.

Underlying Cause. The underlying problem here seems to be that there is no concept of authentication between any parties communicating over the network. As a result, all traffic is trusted and this is clearly a bad transport over which to run critical protocols such as OTAP.

Fix. The obvious action to mitigate this risk would be not to use OTAP. One fix could be to authenticate OTAP messages, e.g. with a scaled down version of the CMS standard for secure firmware update (RFC 4108). Using public key cryptography, generally thought of as computationally infeasible on low-end hardware, might be acceptable for such rare events as authenticating the signature of an OTAP request. Then the network would not be vulnerable to the physical capture of one node and so OTAP could be used securely.

4.7 Remote Command Execution in XServe

Description. XServe is a middleware component that connects the WSN to the back-end. If XServe is started with the `-h` flag, it starts a web server on a user-specified port. This web server is intended to be used to display the output of the attached sensor network. Unfortunately, one of the scripts supplied with the web server contains a bug that can be exploited by an attacker to execute arbitrary commands on the computer running XServe.

Risk. XServe may be run on a Stargate gateway or on the top-tier server. If the web server were enabled on a Stargate and an attacker could access the web server port, she could gain shell level access to the Stargate, obtaining complete control (read, write, modify, drop) over all the data sent from the sensor network to the central servers. If the web server were running on a central database server, an attacker could gain access to the administrative network, potentially gaining complete control over all new data from the network, as well as the ability to modify historical data and attack other parts of the network.

Underlying Cause. This exposure results from a simple programming error, but the underlying problem is really that the XServe component is far too feature rich. From a security viewpoint it is hard to defend the decision of equipping it with a built-in web server, for example.

Fix. We can fix this vulnerability by patching the script or by not using XServe's built in web server. However, since software errors cannot be eliminated, we suggest running complex software on safely stored data, outside the critical parts of the system.

4.8 Stargate Unhardened to IP Attacks

Description. The Stargate we bought new in 2007 shipped with very old versions of several pieces of software, many of which contained exploitable vulnerabilities. Some of the more serious issues are summarised in Figure 2.

Component	Exposure	CVE Reference
OpenSSH	Local root	CVE-2002-0083
		CVE-2003-0682
OpenSSH	Remote root	CVE-2003-0693
		CVE-2003-0695
Kernel	Remote root	CVE-2004-1137
Kernel	Local root	CVE-2005-1263
Kernel	Local root	CVE-2004-1235
PostgreSQL	Remote shell	CVE-2005-0247
PostgreSQL	Remote shell	CVE-2005-0245
PostgreSQL	Remote shell	CVE-2003-0901

Fig. 2. Some vulnerabilities affecting the Stargate. CVE reference numbers may be resolved at `http://cve.mitre.org/cve/`. (Section 4.8).

In addition to these problems caused by outdated software, the Stargate ships in an insecure configuration: it has a default, weak root password and an SSH daemon installed. In addition, if the PostgreSQL database is installed on the Stargate (a recommended and supported configuration), a database super-user is added with a default, weak password. The documentation does not explain how to change these passwords, nor does it suggest that you do. Furthermore the existence of the database super-user account is not even mentioned.

Risks. A remote attacker with access to the Stargate's IP interface may be able to gain root access or crash the Stargate. This would result in either a complete compromise of the sensor network or a complete DoS of the network.

Underlying Cause. The Stargate was not designed for security, as demonstrated by the outdated software, weak passwords, lack of security related documentation and unnecessarily high number of services running.

Fix. Patch, configure, minimise. If possible, update the software on the Stargate, in particular OpenSSH, PostgreSQL and the kernel. Do not run `sshd` on the Stargate and drop all inbound IP traffic to the Stargate that is not already part of an established TCP stream. This implies that all connections to the back-end network would have to be initiated by the gateway. In addition, change passwords for the system root account and the `tele` user in the PostgreSQL system.

5 Architectural Recommendations

Securing the Gateway. The WSN gateway is a bottleneck through which all data related to the sensor network flows. It is therefore crucial that it be properly secured. The usual configuration has the gateway performing many complex operations on the received data and exporting several services to the IP network. We instead recommend not to offer any services to the IP network (e.g. no SSH or HTTP daemons) and not to perform any analysis of the sensor data on the gateway. The gateway should act only as an SSH *client*, rather than server, and simply relay data from the WSN to the server. This would significantly reduce the attack surface of the gateway, resulting in greater security.

Key Management. Clearly, TinySec's default key management approach with a single key for all devices is inadequate to protect an unattended WSN, as the capture of a single node would expose the entire network. The use of pairwise link keys between motes would not protect against attacks on the gateway and another cryptographic protection would have to be used for communication with the central server. Based on the risk assessment in section 3, we recommend the use of pairwise end-to-end keys between each node and the central server to protect the integrity of the sensor data and to provide source authenticity. Insofar as routing is crucial to the security of the network, based on the vulnerabilities exposed in section 4.5 we also recommend the use of pairwise end-to-end keys between each node and the gateway to protect routing table data.

Routing. Many papers on WSNs axiomatically accept the smart-dust-derived assumption that individual nodes (possibly dropped from an aircraft) know nothing about each other's position and that therefore the network can only be built in an ad hoc, decentralised fashion. But, in our scenario, the deployment engineers *know* where they want to place the nodes—for example where they see cracks that need monitoring. We therefore suggest taking advantage of that positional information, perhaps by precomputing an initial (though sub-optimal) set of routing tables and preloading it into the nodes. We also suggest that routing calculations be performed centrally, for example at the gateway, after collecting authenticated local connectivity information from the nodes, since visibility of the whole connectivity graph would allow for the discovery of a better global solution. We are currently working on such a system.

Reviewing the risk assessment. Since the monitoring operation might last for several years, the risk assessment should be repeated at regular intervals to ensure that it still reflects the current usage patterns. It is common for systems to be used in ways that were not originally envisaged, so the protection goals and consequent security measures should be kept up to date.

6 Related Work

Although very many papers have been written about wireless sensor networks, experience papers reporting on real-world deployments are a minority: they include at least Mainwaring et al [6] who monitor seabird nesting environment and

behaviour, Arora et al [7] who deploy a perimeter control WSN and Werner-Allen et al [8] who monitor an active volcano. Closer to our scenario are Krishnamurthy et al [9] who monitor equipment for early signs of failure and especially Kim et al [10] who monitor the structural health of the Golden Gate bridge.

Similarly, although there is a vast literature on security of WSN, including such milestones as Perrig et al [11] on efficient broadcast stream authentication, Eschenauer and Gligor [12] on random key predistribution, Hu et al [13] on secure routing and Chen et al [14] on energy efficient topology maintenance, few such papers deal with attacks on actual sensor network platforms. One notable exception is the brilliant work by Becher et al [3] on physical attacks, providing concrete data on the effort required to extract secrets from a mote. Closer to our own investigations on 802.15.4 jamming are Wood et al [15], who discuss jamming attacks and countermeasures in detail, although their focus is on energy efficiency for the attacker rather than on being able to target a specific victim selectively. The TinySec work of Karlof et al [2] is of particular significance since it resulted in running code which we used extensively. One promising effort in a similar vein is that of Luk et al [16] for which, however, the code was only released after we completed this paper.

Our specific interest lies at the intersection of these two sets: real-world WSN deployments and real-world WSN attacks; so far we have not been able to find other papers in this subset. We believe that practical security is a field worth exploring in greater detail before proceeding with actual deployments.

7 Conclusions

What are the risks of a WSN that monitors large engineering structures? In the situations we examined, the sensors are never linked to actuators; therefore, deploying a WSN in such circumstances does not introduce major new risks. Typically, the worst outcome of an attack is that data gathered from the WSN will be useless, not that the structure itself will be directly endangered. Of course, if the WSN or the collected data becomes useless, there is a financial loss; moreover, false alarms might cause secondary losses through downtime; so, ensuring the integrity of the WSN is still a worthy goal. Considering our qualitative risk assessment (section 3), all the attacks presented in section 4 are potentially relevant, because they can be used directly or indirectly to corrupt integrity of sensor data—the main concern of the structure owners as expressed during interviews. The implementation of countermeasures, which must crucially protect the whole system including back-end and gateway as opposed to just the motes, could be prioritised based on a quantitative risk assessment that established the likelihood of each attack.

Will a WSN built by well-intentioned and security-minded application experts be secure by default? Based on the components we examined, no. For a start, if users of MICAz motes (or any other motes using 802.15.4) wanted to "turn on the crypto", they would find that TinySec does not even compile for their platform. Porting the code is not a trivial endeavour; fortunately, we have

now done it for them. Besides that, we have found and exploited a variety of vulnerabilities. The most devastating in practice are probably the most trivial from an intellectual viewpoint: the Over-The-Air Programming vulnerability (section 4.6), which allows an attacker to reprogram motes at will, and the Remote Command Execution vulnerability (section 4.7), which allows an attacker to gain complete control of the gateway or even the central server. Until such features are patched or protected, it is advisable to keep them disabled.

Among our other attacks, the most significant for the security researcher is probably our sleep deprivation torture attack based on routing table manipulation (section 4.5): it is crippling and very efficient because the attacker can set it up and walk away, as opposed to having to keep talking to the victims to drain their batteries out. Our selective jamming (section 4.3), too, is of interest as it is undetectable by OS and applications and can be used as the basis for more sophisticated attacks including packet rewriting (often assumed possible but rarely demonstrated in a modern ad-hoc radio context) and man-in-the-middle.

We trust that this paper will help vendors and users strengthen the security of their real-world systems.

Acknowledgements

We thank EPSRC for funding this work as part of project EP/D076870/1[10] and the Isaac Newton Trust for co-sponsoring Matt as a UROP summer student. Dan also thanks the MSM 0021630528 project. We are grateful to our industrial partners, interviewed for the risk assessment, and to our colleagues from the Engineering Department of the University of Cambridge for support, especially Neil Hoult who also commented on the paper.

References

1. Kahn, J.M., Katz, R.H., Pister, K.S.J.: Next century challenges: Mobile networking for "smart dust". In: Proc. 5th ACM MobiCom, pp. 271–278. ACM Press, New York (1999)
2. Karlof, C., Sastry, N., Wagner, D.: Tinysec: A link layer security architecture for wireless sensor networks. In: Proc. 2nd SenSys, pp. 162–175 (2004)
3. Becher, A., Benenson, Z., Dornseif, M.: Tampering with motes: Real-world physical attacks on wireless sensor networks. In: Clark, J.A., Paige, R.F., Polack, F.A.C., Brooke, P.J. (eds.) SPC 2006. LNCS, vol. 3934, pp. 104–118. Springer, Heidelberg (2006)
4. Stajano, F., Anderson, R.: The resurrecting duckling: Security issues for ad-hoc wireless networks. In: Malcolm, J.A., Christianson, B., Crispo, B., Roe, M. (eds.) Security Protocols 1999. LNCS, vol. 1796, pp. 172–194. Springer, Heidelberg (2000)
5. Karlof, C., Wagner, D.: Secure routing in wireless sensor networks: Attacks and countermeasures. Ad Hoc Networks 1(2–3), 293–315 (2003)

[10] http://www.winesinfrastructure.org/

6. Mainwaring, A., Culler, D., Polastre, J., Szewczyk, R., Anderson, J.: Wireless sensor networks for habitat monitoring. In: Proc. 1st ACM WSNA, pp. 88–97. ACM Press, New York (2002)
7. Arora, A., Ramnath, R., Ertin, E.: Exscal: Elements of an extreme scale wireless sensor network. In: Proc. 11th IEEE RTCSA, pp. 102–108 (2005)
8. Werner-Allen, G., Lorincz, K., Welsh, M., Marcillo, O., Johnson, J., Ruiz, M., Lees, J.: Deploying a wireless sensor network on an active volcano. IEEE Internet Computing 10(2), 18–25 (2006)
9. Krishnamurthy, L., Adler, R., Buonadonna, P., Chhabra, J., Flanigan, M., Kushalnagar, N., Nachman, L., Yarvis, M.: Design and deployment of industrial sensor networks: Experiences from a semiconductor plant and the north sea. In: Proc. 3rd SenSys, pp. 64–75. ACM Press, New York (2005)
10. Kim, S., Pakzad, S., Culler, D., Demmel, J., Fenves, G., Glaser, S., Turon, M.: Health monitoring of civil infrastructures using wireless sensor networks. In: Proc. 6th IPSN, pp. 254–263. ACM Press, New York (2007)
11. Perrig, A., Szewczyk, R., Tygar, J.D., Wen, V., Culler, D.E.: Spins: Security protocols for sensor networks. Wireless Networks 8(5), 521–534 (2002)
12. Eschenauer, L., Gligor, V.D.: A key-management scheme for distributed sensor networks. In: Proc. 9th ACM CCS, pp. 41–47. ACM Press, New York (2002)
13. Hu, Y.C., Perrig, A., Johnson, D.B.: Ariadne: a secure on-demand routing protocol for ad hoc networks. Wireless Networks 11(1-2), 21–38 (2002)
14. Chen, B., Jamieson, K., Balakrishnan, H., Morris, R.: Span: an energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. wireless networks 8(5), 481–494 (2002)
15. Wood, A.D., Stankovic, J.A., Zhou, G.: Deejam: Defeating energy-efficient jamming in ieee 802.15.4-based wireless networks. In: Proc. 4th IEEE SECON, pp. 60–69. IEEE Computer Society Press, Los Alamitos (2007)
16. Luk, M., Mezzour, G., Perrig, A., Gligor, V.: Minisec: a secure sensor network communication architecture. In: Proc. 6th IPSN, pp. 479–488. ACM Press, New York (2007)

SWF and the Malware Tragedy

Hide and Seek in a Flash

Ben Fuhrmannek[1] and fukami[2]

[1] bef@erlangen.ccc.de, http://pentaphase.de/
[2] SektionEins GmbH, fukami@sektioneins.de, http://sektioneins.de/

December 13, 2008

Abstract The analogy of a children's game of hide and seek describes the constant back and forth of attacks involving Flash on one hand and increasingly sophisticated analysis methods for static code analysis on the other hand. One point in the game utilises heuristic methods to recognise malware, followed closely by simple obfuscation methods to be safely concealed once more.

Introduction

Would you like to play a game of "hide and seek"? Imagine yourself to be just seven or eight years old, playing in the nearby park with your neighbours. First round - you are among the hiding group, scatter out and find a suitable spot behind a huge tree. Soon enough the tree turns out to be a rather obvious choice, too easy to be found. Second round - you are a seeker with a skilled eye for the very hideouts you evaluated a few moments ago for yourself. Since the hiding group knows where you hid in the first place, they try to come up with different places and strategies in order to be concealed for yet another round. Wood is being moved around as diversion or to build new hiding locations. One group disguises as the other to get information on where the others might be hiding or seeking. Successive rounds reveal increasingly complex strategies where each group learns from the other and incorporates previous tactics into their own.

Just like children, virus writers and the antivirus industry have been playing this game for a long time [3]. We can assume, that attackers try to conceal both the existence of their virus code itself and the functionality built into the code. Naturally our seekers in this game try to recognise potential threats and analyse suspicious code.

The major strategy to detect a virus without execution is a signature dictionary: Characteristics of previously classified code form a signature to be used as search pattern. Unfortunately, new or changed code requires a new signature to be detected this way. This weakness is addressed by polymorphic or metamorphic code (see [4]). In order to be still able to detect this constantly changing code, statistical classification comes to mind.

[3] according to [5] the first virus was released in the early 1970s

The same game of hide and seek can be applied to Adobe Flash as well. The seeker's job is to find an attacker's payload before execution. Naturally attackers want to delay the recognition of their payload as long as possible, but at least up to the point of execution. Common attacks involving Flash and static analysis using equivalents to a signature dictionary are described in detail in [2].

Assuming a dynamic payload generation so that static patterns are not applicable as recognition method, it will be shown, how a form of statistical classification can be adapted and how to overcome the heuristics. But first of all, let's have a look around the Flash playground.

Flash Playground

Most commonly an attacker may try to exploit the Flash player, analysis tools, the client (e.g. by stealing credentials) or even try to use a client's CPU time for their own purposes. Although some of these goals may be achieved purely by attacking the format parser of the player or tools, e.g. by fuzzing the SWF format, a parser crash would most likely arise suspicion. Therefore a payload is assumed to be present for the attack.

SWF as defined in [3] is a container format combining arbitrary resources into a single file. Common resources would be images, fonts, audio or video data and some application logic as compiled bytecode. The Flash player contains two entirely separate virtual machines - AVM1 and AVM2[4]. The byte loader feature offered by AVM2, which aims to load data from a byte array during runtime, is of particular interest. This way, encoded, encrypted or otherwise obscured data can be interpreted as SWF without being recognisable as SWF before execution. Starting from Flash version 10 fast memory operations on byte arrays[5] help preparing data to be loaded by the byte loader very efficiently. One example would be an OGG decoder implemented with fast memory operations (see [8]).

Then there is the possibility to program triggers. Arbitrary information available to the Flash runtime environment can be used to trigger the execution of a payload in order to prevent runtime analysis tools detecting an unexpected behaviour. Data such as the Flash player version, browser version, JavaScript interpreter version, IP address, the current date and time and SWF source URL are known to have been used as triggers to effect only specific targets. This data could for example be used to form a crypto key for decoding a payload.

Statistical Classification

For the analysis of non-static code with static function range, a classification method based on n-grams was suggested (see [1]). The front row application for n-grams is language detection of written text. The occurrence of every N successive

[4] AVM1 up to version 8, AVM1 and AVM2 starting from version 9

[5] Efficient fast memory operations[6] have been introduced along with the compiler toolkit Alchemy[7].

characters (including whitespace) of a text is counted and then compared relative to a reference count of known language.

This classification method can be applied to Flash as well. Instead of n characters of a text, a sequence of n opcodes can serve as a basis. For example, a 2-gram profile of the opcode sequence "push, push, add, push, add, return" contains the following pairs of opcode sequences and corresponding counts:

opcodes	count
push, add	2
push, push	1
add, push	1
add, return	1

The distance of two profiles is calculated as simplified distance as suggested in [1]:

$$\delta(SP_r, SP_p) = \sum_{n \in \{op(SP_r) \cap op(SP_p)\}} \left(\frac{2(f_r(n) - f_p(n))}{f_r(n) + f_p(n)} \right)^2$$

where $f_1(n)$ and $f_2(n)$ are the normalised counts of profiles, $op(profile)$ yields all opcodes of a profile and SP_r and SP_p are n-gram profiles.

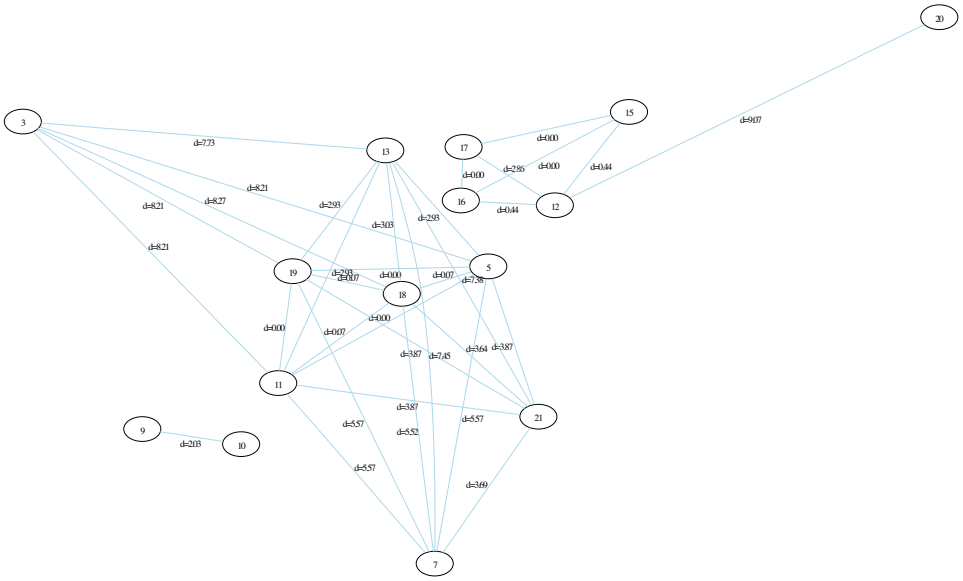

Figure 1. distance graph of n-gram profiles of AVM2 bytecode

Figure 1 shows a graph representation of several SWF9/10 files[6] and their distance based upon the n-gram analysis. Edges connecting far nodes have been eliminated by distance threshold. Three clusters become apparent: {9,10}, {17,16,12,15,20} and {3,13,19,11,18,5,21,7}. Clustering is an expression of similarity between the SWF's bytecode. N-gram profiles contain characteristics of the compiler or IDE, std. libraries, the code's functionality and the code's author(s), each with different intensity. A proof of concept implementation of the Flash adaptation of the algorithm is part of erlswf (see [9]).

Obfuscation and Data Hiding

Initially an attacker would like to conceal the very existence of a payload, however, if found suspicious, it should be as hard as possible to analyse the payload's application logic. These two goals work well together. In order to render the statistical classification ineffective, the payload could either imitate suspected reference characteristics or use as little code as possible. By hiding most of the code, thus leaving only a decoding function exposed to the classification test, both strategies - concealment and obstructing the analysis - can be accommodated.

The following haXe[7] code is a simple example for decoding, loading and executing a payload embedded in SWF using the AVM2 byte loader:

```
class Test {
  static function main() {
    var ldr : Dynamic = new flash.display.Loader();
    var ba : flash.utils.ByteArray = new flash.utils.ByteArray();
    for (i in Data.data) {
      ba.writeByte(i-1); // secret de-obfuscation algorithm: i - 1
    }
    ldr.loadBytes(ba);
    flash.Lib.current.addChild(ldr);
  }
}
```

In a similar way, a payload could have been loaded from an external source, e.g. a picture from the internet, decoded, possibly decrypted, then loaded and executed.

Conclusion

Statistical classification can be effectively applied to identify bytecode similar to a reference profile only up to the point, where characteristics of the reference profile can be approximated by a payload or too few data exists to analyse

[6] files were chosen to contain a suitable amount of AVM2 bytecode
[7] see [10]

bytecode effectively. In the end it always breaks down to a matter of how much effort each side is willing to put into either obfuscation or analysis respectively. Malware seen in real-life at the moment hardly makes any effort to be not easily detected by heuristic methods. This is expected to change in the future, as well as more use of the byte loader and fast memory operations.

Also consider the possibility of one side of the hide and seek game to be tricked into making the next move, just to be observed by the opponent and learn from the reaction. What if you don't even know whether your are currently playing the game or not? You may be hiding accidentally right now and don't even realise it. Every new pattern recognition algorithm, coding algorithm or obfuscation idea will be taken into account and can most likely be adapted to serve as the next round of the game. So the question remains: Would you like to play a game of "hide and seek"?[8]

References

1. J. Filipe et al. (Eds.): Supprting the Cybercrime Investigation Process: Effective Discrimination of Source Code Authors Based on Byte-Level Information - ICETE 2005, CCIS 3, pp. 163173, Springer-Verlag (2007)
2. fukami and Ben Fuhrmannek: SWF and the Malware Tragedy: Detecting Malicious Adobe Flash Files, https://www.flashsec.org/mediawiki/images/5/57/SWF_and_the_Malware_Tragedy.pdf, March 9, 2008
3. Adobe Inc., SWF Format Specification version 10, Specification, http://www.adobe.com/devnet/swf/
4. Wikipedia: The Free Encyclopedia: Polymorphic Code, http://en.wikipedia.org/wiki/Polymorphic_code
5. Wikipedia: The Free Encyclopedia: Computer Virus, http://en.wikipedia.org/wiki/Computer_virus
6. Scott Petersen, FlaCC, Presentation Slides, http://llvm.org/devmtg/2008-08/Petersen_FlashCCompiler.pdf, August 1, 2008
7. Adobe Alchemy, Tool, http://labs.adobe.com/technologies/alchemy/
8. Adobe Alchemy, Code Example, http://labs.adobe.com/wiki/index.php/Alchemy:Libraries
9. Ben Fuhrmannek, erlswf: Toolkit for disassembling SWF up to version 10. Tool, http://code.google.com/p/erlswf/
10. Motion-Twin, haXe, Tool, http://haxe.org/
11. Ben Fuhrmannek and fukami: SWF and the Malware Tragedy, 25C3 Presentation, http://events.ccc.de/congress/2008/Fahrplan/events/2596.en.html, December 2008

License

[8] For more elaborate answers to this question see [11].

The Ultimate C64 Overview

Michael Steil, http://www.pagetable.com/
25th Chaos Communication Congress 2008

Retrocomputing is cool as never before. People play C64 games in emulators and listen to SID music, but few people know much about the C64 architecture and its limitations, and what programming was like back then. This paper attempts to give a comprehensive overview of the Commodore 64, including its internals and quirks, making the point that classic computer systems aren't all that hard to understand - and that programmers today should be more aware of the art that programming once used to be.

Commodore History

Commodore Business Machines was founded in 1962 by Jack Tramiel. The company specialized on electronic calculators, and in 1976, Commodore bought the chip manufacturer MOS Technology and decided to have Chuck Peddle from MOS evolve their KIM-1 computer kit (a design that demos their new MOS 6502 8 bit CPU) into the Commodore PET series: computers with built-in monitors for the home, school and small business market that ended up competing with devices from Atari and Apple.

In 1981, Commodore introduced the VIC-20, a 5 KB stripped down monitorless computer-in-the-keyboard design based on the PET for the home computer market. This was followed by the (incompatible) higher-end Commodore 64 in 1982 that included more PET features, came with 64 KB of RAM (an immense amount compared to the rest of the market) and was very aggressively priced at US$595 beating the competition by a factor of two. This was made possible by designing and building most of the system in-house.

Although some features of the C64 were taken from the PET models of the time, it had to be connected to a TV set, which only made 40 columns of text possible (as opposed to 80 columns on the PET). Also, the BASIC 4.0 codebase was stripped down to the old 2.0 feature set to make it fit into 8 KB.

In the beginning, the C64 did well in the competition. The superior but compatible C128 from 1985 did well, too, but was never more popular than the C64, which continued to be sold. The direct successor to the low-end VIC-20, the 1984 Plus/4 and its siblings, the C16 and the C116, failed, mostly because they were incompatible with the C64, which at that time already had a remarkable software library.

A few years after the introduction, the C64 was still offered as a low-end alternative to the Commodore Amiga, and while it became less popular in the USA, it gained more and more popularity in Europe. In the early 90s, when manufacturing costs of a C64 were as low as $25, it gained a second life in Eastern Europe. Production did not end until the liquidation of Commodore in 1994. According to the 1993 Annual Report, 17 million C64 had been produced in by then, as well as 4.5 million C128.

Look and Feel

A C64 only needs to be connected to power and a TV set (or monitor) to be fully functional. When turned on, it shows a blue-on-blue theme with a startup message and drops into a BASIC interpreter derived from Microsoft BASIC. In order to load and save BASIC programs or use third party software, the C64 requires mass storage - either a “datasette” cassette tape drive or a disk drive like the 5.25" Commodore 1541.

Unless the user really wanted to interact with the BASIC interpreter, he would typically only use the BASIC instructions LOAD, LIST and RUN in order to access mass storage. LOAD"$",8 followed by LIST shows the directory of the disk in the drive, and LOAD"filename",8 followed by RUN would load and start a program. If a tape drive is connected, LOAD and RUN will launch the first program on tape - pressing SHIFT and the RUN/STOP key has the same effect.

By default, typing characters without SHIFT will result in upper case characters being shown on the screen. This can be changed by pressing the Commodore key and SHIFT at the same time, which switches to the upper/lower character set. This behavior is due to the fact that the first PET only had uppercase characters, and that BASIC required keywords unshifted - but all tutorials taught them as uppercase. This is also the reason why in Commodore's version of ASCII, called PETSCII, the codes for uppercase and lowercase characters are reversed.

The text based user interface is not a line editor, but a screen editor, i.e. the cursor can be moved freely on the screen, and pressing RETURN on a line with existing text will present the text to the application as if it were just typed in.

While a physical screen line is only 40 characters, the screen editor logic can extend it to a logical 80 characters - whenever a character is entered on the 40th column, the rest of the screen is moved down by one line, ”opening” a line that extends the previous line to 80 characters.

The C64 screen editor supports selecting one out of 16 color for the foreground text by pressing the key combinations Ctrl+1 to Ctrl+8 and Commodore+1 to Commodore+8. Reverse text mode can be turned on and off with Ctrl+9 and Ctrl+0.

Ports and Connections

The C64 has a whole range of connection possibilities. On the side, it has two 9 pin Atari-style joystick connectors that can also be used for a mouse, light pens or paddles. On the back, there is the expansion port, which exports the complete processor bus, allowing not only game cartridges but also cartridges with I/O chips that map themselves into the CPU's address space - or even cartridges that completely replace the CPU. For a conventional TV connection, there is an

RCA connector that outputs an RF signal. For monitors, there is an extra DIN connector that carries separate chroma, luma and audio signals (S-Video).

For connecting Commodore compatible printers and disk drives, there is a DIN connector for the IEC bus. There is also a dedicated connector for datasette drives. The ”User Port” consists of several GPIO pins that can be used for custom hardware projects, or as a RS-232 port (with TTL levels), for which support exists in the ROM.

Board

On the C64 motherboard, there is a dedicated IC each for the main tasks. There is the MOS 6510 CPU, eight 64 KBit RAM chips (later consolidated into 2), three ROM chips with KERNAL (I/O library), BASIC and the character set (KERNAL and BASIC were later consolidated), two 6526 CIA I/O controllers (one for keyboard and joystick, one for the IEC bus and the user port), the 6581 SID sound chip, and the 6567/6569 VIC video chip, as well as the RAM chip that holds the 512 bytes of Color RAM.

All non-RAM chips are custom chips designed manufactured by MOS, Commodore's inhouse chip company.

Address Space

The 8 bit C64 design has a 16 bit address bus, allowing the CPU to address 64 KB of memory. Since the C64 has 64 KB of RAM, filling the complete address space, ROM and I/O chips are mapped into regions of the address space that are shared with RAM: The CPU can switch these regions between RAM and a second or third mapping. These regions are as follows:

- $0000-$9FFF: RAM
- $A000-$BFFF: RAM or BASIC ROM
- $C000-$CFFF: RAM
- $D000-$DFFF: RAM or memory mapped I/O chips or character ROM
- $E000-$FFFF: RAM or KERNAL ROM

In contrast to CPUs like the Z80 and the 8086, and like most modern CPUs, I/O devices are memory mapped on the C64's 6510 CPU. The mapping is as follows:

- $D000-$D3FF: VIC video controller
- $D400-$D7FF: SID sound controller
- $D800-$DBFF: Color RAM
- $DC00-$DCFF: CIA 1 I/O controller
- $DD00-$DDFF: CIA 2 I/O controller
- $DE00-$DFFF: for extensions on the expansion port

6502 CPU

The CPU inside the C64 is a 0.985 MHz (on PAL) MOS 6510, which is a close derivative of the well-known 8 bit little-endian MOS 6502. The 6502 was introduced in 1975 by MOS Technology, a company formed earlier the same year by former Motorola engineers, headed by Chuck Peddle. The philosophy of the 6502 was to have a reduced instruction set and a small register file, making it simpler and faster than CPUs like the Z80 at the same clock speed, as well as cheaper to manufacture.

Unlike other CPUs, the 6502 does not have a set of general purpose registers. Instead, it has a single accumulator A (for arithmetic and logic), two index registers X and Y (for incrementing, decrementing and indexing memory) and a stack pointer. All these registers are 8 bits. The processor status, consisting of the negative (N), overflow (V), break (B), decimal (D), interrupt (I), zero (Z) and carry (C) flags is exposed as register P. The program counter (PC) is 16 bits wide. The fact that the stack pointer is 8 bit means that the stack is confined to the area between $0100 and $01FF in the address space, i.e. the upper half of the effective stack pointer is hard-coded to $01. There is another special area in the address space: The first 256 bytes, at $0000 to $00FF are referred to as the zero page (ZP). Many instructions support special encodings for zero page addresses, which saves one byte in the instruction encoding as well as at least one cycle of execution time. This can be seen as an extension of the register file to another 256 (though external) registers.

All opcodes are one byte, and have 0, 1 or 2 byte operands. The 8x8 opcode matrix is somewhat logical (e.g. branch instructions are encoded as $10, $30, $50, ...), but there is no easy rule to construct the opcode table. Nevertheless, the opcode table is a minimal encoding for optimal decoding in the 6502's internal PLA ROM.

Instruction Set

The instruction set is very streamlined, and avoids redundancies. There are load instructions (LDA/LDX/LDY to load A, X and Y respecively), store instructions (STA/STX/STY), read-modify-write instructions (logic: ASL/LSR/ROL/ROR, count: INC/DEC), arithmetic (ADC/SBC; note that these always include the carry: CLC/ADC is a regular addition, and SEC/SBC is a regular subtraction, because of the one's complement logic), compare (CMP/CPX/CPY; these are subtractions without storing the result), logic (AND/ORA/EOR, and BIT, which is AND without storing the result), as well as branch instructions, flag manipulation, register transfer and stack manipulation.

Addressing Modes

Each instruction supports one or more addressing modes. Common instructions like LDA (load accumulator) support more addressing modes than less common ones (BIT).

- The immediate addressing mode is indicated with a # sign: LDA #$17 loads the immediate value of $17 into the accumulator.
- Absolute addressing specifies a 16 bit address as an operand: LDA $0314 loads from the memory address $0314.
- Zero page addressing is an optimized version of absolute addressing: LDA $02 will read from ad-

dress $0002 in memory, but the instruction can be encoded more tightly, and execution is faster.

- Absolute-X-indexed addressing reads from a specified address, to which the contents of the X register is added. LDA $0200,X reads from the address $020A, in case X is $0A. This allows reading from tables.
- Absolute-Y-indexed is the same thing, but with the Y register.
- Zero-Page-X-indexed is an optimized version of Absolute-X-indexed. LDA $F0,X reads from the Xth location in a table stored starting at $00F0 in memory. Note that zero page addresses will wrap around, so $F0 + $10 = $00.
- Zero-Page-Y-indexed is the same thing, but with the Y register.
- Zero-Page-X-indexed-indirect adds X to a specified zero page address, reads a 16 bit pointer from the resulting address and finally accesses memory at that address. So LDA ($80,X) will read from an address specified by the array of pointers at $0080 and the index X into the array. This addressing mode is rarely used.
- Zero-Page-indirect-Y-indexed treats two consecutive bytes in zero page as an address and adds Y to the address. LDA ($14),Y will read from $E020, if the address stored at $14/$15 is $E000 and Y is $20. This addressing mode is the most convenient way to work with pointers, as no register can hold 16 bits.

Register Transfer and Stack

There are several 1 byte instructions without operands that move data between registers. TAX, TXA, TAY and TYA move between A, X and Y. TSX and TXS copies between X and the stack pointer.

The stack pointer always points to the next address that is written to. This means that an empty stack has a stack pointer of $FF, and pushing a value first writes the value and then decrements the stack pointer. The 6502 can move the accumulator from and to the stack (PHA/PLA), as well as the processor status P (PHP/PLP).

Control Transfer

Next to the absolute JMP instruction, there is an indirect version that jumps over a vector (e.g. JMP ($FFFC)). JSR (jump to subroutine) only has an absolute version, and stores the address of the next instruction minus one on the stack. RTS (return from subroutine) takes the address from the stack, adds one, and moves it into the program counter. The ”minus one” logic was chosen because it could save one cycle in the implementation of JSR.

A hardware interrupt, unless disabled by a set interrupt (I) flag, pushes the address of the next instruction minus one (just like JSR), pushes the processor status afterwards, disables interrupts, and jumps over the vector at $FFFE/$FFFF. RTI return from interrupt) is the same as the combination of PLP and RTS. BRK causes a software interrupt and behaves the same as a hardware interrupt, except that it sets the B flag on the stack to 1 (a hardware interrupt sets it to 0).

NMIs behave the same as IRQs, but they cannot be masked, and they use the $FFFA/$FFFB vector. The reset vector is at $FFFC/$FFFD.

Flags and Branches

All load and logic instructions set N and Z accordingly, shift instructions also modify C, and arithmetic instructions touch N V, Z and C. The D (decimal), I (interrupt disable) and C flags can be set and cleared programmatically (CLD/SED, CLI/SEI, CLC/SEC), while the V flag can only be cleared (CLV). Conditional branches are possible based on the value of the negative (BPL/BMI), overflow (BVC/BVS), zero (BNE/BEQ) and carry (BCC/BCS) flags. Branches encode an 8 bit relative offset and can therefore reach code in the area of +127 and -128. Since a compare is the same as a subtraction, BCC is a branch on (unsigned) below, and BCS is a branch on above-or-equal.

NOP

NOP (no operation) does nothing. Its encoding is $EA.

Decimal Mode

If the D flag is set, all ADC and SBC operations will be BCD-adjusted afterwards, i.e. $09+$02 won't be $0B, but $11, since 9+2=11. The BCD correction circuit has been patented in US patent 3,991,307.

Cycle Counting

It is quite straightforward to find out how many cycles an instruction takes. As a rule of thumb, an instruction takes as many cycles as the number of memory fetches it has to perform, but at least two.

Therefore, single-byte opcodes (one byte fetches; NOP/TAX/INX etc.) as well as instructions with immediate operands take two cycles. Zero page instructions take three memory accesses (opcode, address, data), so they are three cycles. Absolute instructions take four accesses (opcode, address low, address high, data), so they are four cycles.

Read-modify-write instructions (INC/DEC/shift/rotate) are an exception: They require 4 memory accesses for the zero page case and 5 otherwise, but they take 5 and 6 cycles, respectively.

Branches take 3 cycles if they are taken and two if they are not taken. And extra cycle has to be added if the branch crosses a page boundary. JMP is 3, push is 3, pull is 4, JSR and RTS are 6 each.

All other timings can be looked up in the 6502's reference, but they are very easy to memorize.

Common Tricks

- The BIT instruction exists in a two-byte (immediate operand) and three-byte (absolute operand) variant. Since BIT only changes the flags, it effectively skips one or two bytes in the instruction stream. This can be used to replace a two-byte branch or a three-byte JMP with a one-byte BIT if only one or two bytes have to be skipped.
- The architecture allows safe self-modifying code, so a common optimization for copy loops is to use LDA $nn00,X and STA $mm00,X, looping X from $00 to $FF and then incrementing the bytes that encode nn and mm for the next page. Compared to

a LDA (zp1),Y: STA (zp2),Y sequence, this gets the inner loop down from 16 cycles (5 LDA, 6 STA, 2 INY, 3 BNE) to 14 (4 LDA, 5 STA, 2 INY, 3 BNE).

- A PHA/PLA combination is 7 cycles, while an STA/LDA combination in the zero page is 6 cycles, so unless there is no free zero page space, PHA/PLA should be avoided to quickly store a value. Using an absolute store to write the value into the operand of a future immediate load (i.e. self modification) is the same speed at the zero page solution, but does not waste zero page space.
- An elegant way to store a flag is to have it in bit #7 of a zero page address. While a load/store combination has to be used to set the flag, it can be cleared with a simple LSR (5 cycles) and tested with BIT (3 cycles), without affecting register contents.
- Since the 6502 is so register starved, only 3 bytes can be passed to a subroutine in registers. Also, the stack is small, and accessing it is slow, so stack frames as seen on modern architectures are very uncommon. Many applications and libraries (e.g. GEOS) use a dedicated area in the zero page as virtual registers.
- The 6502 has no instructions for multiplication, division or floating point arithmetic. Most 6502-based computers have a BASIC interpreter in ROM though, and they typically include a math and floating point library.

Bugs and Quirks

The original 6502 implementation has a series of bugs and other anomalies that have never been fixed in MOS chips (not counting the 65CE02, which was only used in Amiga peripherals).

- The indirect version of JMP loads the program counter from the wrong address, if the vector's address lies on a page boundary: JMP ($23FF) will read the address from $23FF and $2300 instead of $23FF and $2400.
- When in decimal mode, the negative flag reflects the original binary result, not the effective decimal result.
- If a software interrupt (BRK) and a hardware interrupt occur at the same time, the BRK is dropped.
- Read instructions (LDA/AND etc.) with the absolute-indexed addressing mode first read from the absolute address without the index register added, and then read again from the correct address. LDX #$07 LDA $D019,X will first read from $D019, discard the result, and then read from $D020. On the C64, this read form $D019 would ACK all pending VIC interrupts, while it is only supposed to read the border color ($D020).
- Read-modify-write instructions with absolute addresses first read the value, but one cycle before they store the result, they store the original value again. On a C64, this can be seen when incrementing the screen border color at a defined area of the screen, as every write to the register will cause a tiny gray dot on the screen. When this is used with certain I/O ports, this can have other side effects. The latter two quirks have been used heavily for obfuscating copy protection software.
- Instruction decoding in the 6502 is done by a PLA that compares the current cycle number within the instruction and the current opcode against a ROM of 130 mask lines, of which any number can fire independently. The outputs of these lines are then fed into components like the ALU, bus control, register control and program counter logic. The instruction set only consists of 151 defined opcodes, and since handling the remaining 105 opcodes as NOPs or traps would have required extra lines in the PLA, they will match against some lines that were meant for instructions with similar opcodes. Some of these ”illegal opcodes” lead to useful results and are used in some software (SAX = store A & X), but most of the instructions make little sense (SHX = store X & the upper 8 bits of the program counter), and some even lock up the CPU, disabling IRQs and NMIs (CRA/KIL).

The MOS 6510

Except for the pin layout, the MOS 6510 that is used in the C64 differs from the generic MOS 6502 in two ways: It can make the bus tri-state when not used, so the VIC can use it, and it has a 6 bit I/O port built in, which can be controlled using zero page locations 0 and 1. In register 0, each bit from 0-5 set it to output if 1, and to input if 0. Bits 0-5 in register 1 are the actual I/O pins. On the C64, bits 0-2 are outputs and control bank switching, they turn the ROMs and the I/O area on and off. Bits 3-5 go to the tape connector and control the motor and the data sent to the head, and detect whether a key on the tape deck is pressed.

BASIC

Microsoft had a strong position in the market for (mostly ROM) BASIC interpreters in 8080-based home computers when the MOS 6502 was released in 1975, so they rewrote their interpreter in 6502 assembly. Microsoft BASIC was pure 6502 code with a minimal character I/O interface to the machine's ”monitor”, i.e. I/O library.

Commodore decided to license the interpreter for the 1977 PET and extended it slightly to interface with their disk and tape libraries. Commodore BASIC was very buggy, so they went back to Microsoft for an update, which, with the Commodore changes re-applied, shipped in newer PETs as BASIC V2. For version 4, Commodore added several extensions, both language constructs as well as support for graphics and sound.

Being a low-end machine, Commodore took the bug-fixed BASIC V4 codebase and removed all features after V2, making it independent of the machine's graphics and sound features again, and fitting it back into 8 KB, and shipped this version on the VIC-20. The Commodore 64 got the exact same version, except that it runs at a different memory address ($A000-$BFFF).

Microsoft BASIC is a line-based editor, that is, lines can be shown with the LIST command, and they can be modified by re-typing them. This integrates nicely with the KERNAL screen editor: The cursor can be moved up to LISTed lines, the lines can be modified, and when RETURN is pressed, the whole line is fed into BASIC again.

A nice feature of this and later versions of Commodore BASIC is the fact that all important parts, like the tokenizer, the detokenizer and the interpreter loop jump over a jump table in RAM before they do their work, allowing the user to extend BASIC arbitrarily. The most well-known BASIC extension is Simons' BASIC, a cartridge that maps 8 KB of extra ROM at $8000-$9FFF.

KERNAL

The C64 has an 8 KB I/O library at $E000-$FFFF which is utilized by BASIC, but is intended to be used by other applications as well. All Commodore 8 bit systems have a standardized library call interface in the form of jump tables at the very top of memory that call into machine-specific functions for I/O.

KERNAL is started form the RESET vector, initializes the machine, sets up an interrupt service routine that handles the keyboard, animates the cursor and does the real time clock. The C64 has a hardware clock in each of the CIA chips, but KERNAL has not been updated to use this feature since the VIC-20.

KERNAL provides an abstract character I/O interface to a number of devices. All devices support open, read, write and close. The open call takes three parameters: The logical file number (there is a maximum of 16 channels), the device number and the secondary address. There are 16 device numbers statically assigned to the devices. The 8 bit secondary address can signal something to the device, like speed or an operation mode. Some devices (tape and IEC) support an optional filename.

Device 0 is the keyboard. While KERNAL exports raw key presses, the keyboard can also be accessed through character I/O, which will go through the screen editor and replay all characters on the screen that are in the line of the cursor, regardless of whether the user typed them or they had been there before.

Device 1 is the tape drive. KERNAL reads and writes blocks of data at a time and buffers them for character I/O.

Device 2 is RS-232. KERNAL contains a very sophisticated (but rarely used) software RS-232 implementation that supports up to 2400 baud.

Device 3 is the screen. KERNAL interprets special codes, manages the cursor position and handles scrolling.

Devices 4 to 15 will be directed to the IEC bus. By convention, devices 4 to 7 are printers and plotters, and devices 8 to 15 are floppy or hard disks.

KERNAL allows interacting with the IEC bus manually by sending TALK and LISTEN requests to the bus.

GEOS

While KERNAL is a minimal character-based operating system in ROM, there is also a disk-based operating system with a graphical user interface for the C64. GEOS was released by Berkeley Softworks in 1986 and Commodore bundled it with the C64 for some time. The GUI, which can be controlled by a joystick or a mouse, runs in 320x200 graphics mode and resembles early versions of MacOS. GEOS is a 16 KB library that includes an optimized disk interface (faster, support for timstamps, icons and multi-fork ”VLIR” files), library code for drawing to the screen, high level UI primitives for menus, buttons and dialog boxes (with callbacks) and a simple memory swapping facility. Furthermore GEOS allows input and printer driver plugins, as well as proportional fonts in different sizes. Internally, GEOS has a jump table to its library routines that consists of about 150 entries.

GEOS came with applications like GeoPaint and GeoWrite; Berkeley Softworks themselves offered solutions like GeoPublish and GeoCalc, and more software was available from third parties.

GEOS' only requirement is a 1541 disk drive, but a 3.5" 1581 drive, a RAM extension or one of the later hard drives helped speed it up a lot.

6526 CIA

The C64 has two identical 6526 CIAs (Complex Interface Adapter) that are mostly used for I/O. One CIA features 16 general purpose I/O pins (8 bit port A and 8 bit port B) that can be used either as an input or an output, two programmable timers and a real-time-clock.

The timers have 16 bit counters and count down by one either on each clock cycle, or on an external event, or on a timer A underflow (in the case of timer B). This allows concatenating the timers to one 32 bit timer. On an underflow, the CIA can be programmed to cause an IRQ or to send data through a serial shift register. The CIA also supports receiving data through a shift register.

The real-time-clock has a resolution of 1/10 of seconds and supports generating interrupts at a certain time.

CIA 1 is hooked up to the keyboard and the joystick ports. Since the keyboard consists of 64 keys (plus SHIFT LOCK, which is parallel to the left SHIFT key, and RESTORE, which is directly connected to the CPU's NMI line), these can be laid out in a 8x8 matrix of lines, key presses connecting the intersections. One side of the matrix is connected to port A (output), and the perpendicular side is connected to port B (input). The keyboard driver can now write the values of $01, $02, $04 etc. in port A and test the input of port B to see which keys are pressed. The two joysticks are connected in parallel to port A and port B, so they can cause spurious keyboard events.

CIA 2 is hooked up to the IEC bus, and I/O lines control the VIC bank. The rest is exposed on the user port, and can be used for RS-232.

KERNAL uses CIA 1 for the 50 Hz system timer, but, apart from the ports, doesn't use any of the extra features of either CIA.

6581 SID

The SID (Sound Interface Device) is a whole topic of its own.

6567/6569 VIC

The video chip inside the C64 is called the MOS 6567/6569 VIC-II (Video Interface Controller) - the video chip in the VIC-20 had been the original VIC, and was the reason for the marketing name of the VIC-20.

The VIC supports a 40x25 text mode, a 320x200 bitmap mode, 16 colors and 8 sprites - all of these features have lots of sub-modes and options. The amount of memory the VIC can address is 16 KB, and while by default, it accesses the first 16 KB of the C64 RAM, it can be configured to use any of the four banks.

The 16 colors of the VIC are divided into two sets of eight. The first eight are the more important colors, as some modes only support the first eight. The colors are, in the original order: black, white, red, cyan, purple, green, blue, yellow, orange, brown, pink, dark gray, gray, light green, light blue and light gray.

Character Mode

The C64 has two built-in character sets that the VIC can access. They can be shown on the screen by writing the numbers 0 to 255 into screen RAM (at $0400 by default). The default font has uppercase characters and lots of line-drawing symbols in the lower 128 characters, and the second half consists of the same characters, but inverse. The alternative font has upper- and lowercase characters and omits some of the symbols.

The foreground color of the characters can be changed by writing the color numbers into the Color RAM, which is located at $D800-$DBFF. There is one byte per character, but only the lower 4 bits are actually preserved by Color RAM.

Each character is 8x8 pixels, and stored as eight bytes in the character ROM, one line being one byte. A 1-bit will take the color from the Color RAM ($D800-$DBFF), and a 0 bit will take it from the global background color register ($D021). The pixel matrix is determined by looking up the character index in the screen RAM (at $0400-$07FF by default) and consequently looking up the pattern the current character set (the VIC sees the default font at $1000, although for the 6502 it is invisible there).

In Extended Color Mode (ECM), it is possible to chose between one of four background colors (registers $D021 to $D024) with the upper two bits of the character index, but then only 6 bits will be used to look up the character pattern, decreasing the number of possible characters to 64. The built-in (uppercase) character set is well-suited for this: While it is similar to the ASCII encoding, it has the uppercase characters mapped to codes $01 to $1A, so the most important characters are within the first $40.

Multi-Color Character Mode allows up to four colors per character and is intended for tile-based games, like platformers. If bit 3 of the value in Color RAM is 0, then the character gets displayed just like in non-multicolor mode, but colors are restricted to the first eight. If bit 3 is 1, then pairs of horizontally adjacent bits are combined in their meaning: 00 represents the screen background ($D021), 01 is the second background register ($D022), 10 is the third background register ($D023), and 11 is the color specified in bits 0-2 of the Color RAM. Pixels in these characters are twice as wide, so the resolution of a character is 4x8.

Bitmap Mode

In hi-res graphics mode, the VIC supports a resolution of 320x200, which uses the same pixel frequency as 40x25 character mode (40*8=320, 25*8=200). The bitmap can reside at $0000 or $2000, and the VIC reads one byte for 8 pixels. But hi-res mode does not only support monochrome graphics: The foreground and background colors of each 8x8 tile are taken from the high and low nibble of screen RAM, which would otherwise be unused. Color RAM is not used in this mode.

The encoding of the bitmap is identical to the encoding of a character set, making it a non-linear framebuffer: The first eight bytes of the bitmap represent the pixels in the tile at character position (0,0), the second eight bytes represent the tile at character position (1,0), which is pixel position (8,0), and so on. This layout makes pixel addressing in software slower.

In Multi-Color Bitmap Mode, the horizontal resolution is halved to 160x200, and pixels are twice the width. Every set of two bits in the encodes one of four colors per tile: 00 takes it from the global background register ($D021), 01 and 10 take it from the upper and lower nibble of screen RAM, respectively, and 11 takes it from Color RAM.

Scrolling

The VIC supports hardware X and Y scrolling by 0 to 7 pixels. Since the 40th column is half visible and another column left of the first column is half-visible when the horizontal shift register is set to e.g. 4, a 41th column would be needed. Instead, it is possible to switch the screen to 38 column mode, i.e. the whole screen is a little narrower, and more border is shown on the left and on the right. The screen can also be switched from 25 to 24 lines the same way.

Sprites

The VIC has eight hardware sprites (also called MOBs, movable objects). Each sprite is 24x21 pixels, which is encoded in 63 bytes. Set bits will be drawn in the sprite's individual foreground color, and cleared bits will be transparent. The index to the sprite's bitmap data in memory is an 8 bit value that is read from the last 8 bytes of screen RAM - since the screen is only 1000 characters, the last 24 characters of the $0400=1024 bytes area would otherwise be unused.

There is also a multicolor mode for sprites, which makes pixels twice as wide and decreases the horizontal resolution to 12 pixels. In this mode, 00 is still transparent, and 10 encodes the sprite's individual color. The codes 01 and 11 take the color out of the sprite multicolor registers ($D025 and $D026), which are shared among all sprites.

Sprites can be positioned at arbitrary pixel positions on the screen, and overlap. In this case, sprites with lower numbers have priority over sprites with higher numbers. Each sprite can either be shown in front or behind background pixels. Sprites can be X- and Y-expanded by a factor of two, and collision of two sprites or of a sprite and background pixels is de-

tected by hardware: Whenever two non-transparent sprite pixels are drawn at the same position on the screen, they have collided. Whenever a non-background pixel is drawn by the character generator at the same position where a non-background pixel of a sprite is draw, the sprite has collided with the background. An exception from this rule is the 01 code in the character data, which also counts as background. This way, a background picture can be constructed that does not cause collisions in certain areas. In practice, most newer games do not use the hardware functionality, but instead test for overlapping sprite bounding boxes in software.

Memory Layout

The VIC can address 16 KB at a time. All VIC data structures can be stored anywhere in these 16 KB, but they have to be aligned to their size.

- The screen RAM is $0400 bytes in size and can be at $0000, $0400, ...
- The character set is $0800 bytes, and can be at $0000, $0800, ...
- The bitmap is $2000 bytes and can be at $0000 or $2000.
- Sprites are $40 bytes, and can be at $0000, $0040, ...

Two GPIO pins of the second 6526 CIA are connected to bits 6 and 7 of RAM when the VIC accesses it. By changing the lower 2 bits of $DD02, the VIC can be switched between banks $0000-$3FFF (11), $4000-$7FFF (10), $8000-$BFFF (01) and $C000-$FFFF (00).

If the VIC is set to banks $0000 or $8000, then the two built-in character sets shadow RAM in the area of $1000-$1FFF. This means that the built-in character set can be used on those banks without occupying RAM, but it also means that the area from $1000-$1FFF cannot be used for bitmap, screen RAM or sprite data either.

For timing reasons, color information is not taken from main RAM, but from a dedicated Color RAM. These $0400 half-bytes are accessible to the C64 at $D800-$DBFF and can not be bank switched.

Timing

For advanced VIC programming, it is necessary to not just set up a certain mode and have the VIC display it, but to reprogram the VIC while it is drawing the picture. For this, it is necessary, to understand its timing.

While the pixels within the screen area are 320 by 200, the VIC actively draws pixels in the border color outside of this area, which (on PAL) is 403x284 pixels. Analog TV standards specify an H blank area at the end of every line, and V blank area at the end of every screen. So counting this timing as pixels, this gives an absolute resolution of 504x312 pixels. The interesting and very useful connection about the pixel clock and the system clocks is that an 8 pixel character is drawn every system clock cycle, i.e. about 1 million times a second. The 504 horizontal pixels therefore mean that a line is drawn on the screen every 63 cycles. With this information, it is possible to do cycle-exact timing of assembly code to switch a VIC register at an 8 pixel granularity.

Further timing details (badlines, sprite timing), as well as the application of this information to do tricks like FLD, FLI and AGSP would go beyond the scope of this article, but are talked about in the presentation at the 25th Chaos Communication Congress.

Memory Configuration

In a running system with BASIC and KERNAL, the BASIC and KERNAL ROMs are turned on and visible at $A000-$BFFF and $E000-$FFFF respectively, and at $D000-$DFFF, the I/O area is visible. Using the 3 lowest bits in the processor port at address 1, this configuration can be altered. The ROMs can be turned off, revealing RAM instead, and the I/O area can be configured to show either RAM or the character set ROM. Note that writing to ROM will always direct the data to the underlying RAM.

In practice, many programs run in the default configuration and use both KERNAL library routines, as well as functions in BASIC, to keep their own code as small as possible. More recent programs and almost all games turn off all of ROM, to get direct control of the interrupt vector without having to go though the KERNAL handler first. The I/O area is typically configured to show the I/O registers and Color RAM, and only rarely switched to a different configuration to temporarily access the RAM underneath. A few applications read the character ROM at program start and modify the copy.

The C64 supports another ROM bank at $8000-$9FFF, which can only be serviced by an external cartridge connected to the expansion port. If KERNAL detects the magic string ”CBM80” at $8004 on startup, it will jump to the code of the cartridge right away.

Tape Interface

The tape interface consists of a single line each for data input and output, motor control and key sense. The raw data is read from and written to the data lines, and all encoding and decoding of the data stream is done in software. 3 of the required lines are connected to the processor port at zero page location 1, and one (data input) is connected to CIA 2.

IEC Bus

The IEC bus is a serial version of the IEEE-488 bus used on the PET. Devices on the IEC bus are daisy-chained, and are all connected to the same three lines: ATN (attention), clock and data. IEC has a single bus master, which is the computer. It is the only device to ever raise ATN, while every device can output to clock and data, depending on the state of the bus.

If the computer raises ATN, every device on the bus listens for the 4 bit device number and compares it with its own. The protocol on who sends and who receives is determined by the computer sending TALK/UNTALK and LISTEN/UNLISTEN requests, quick is

an ATN sequence, followed by one of the four commands.

While KERNAL exports the interface at this level, it also allows high-level open, close, read and write operations on the IEC bus, as well as load and save operations. The BASIC LOAD and SAVE commands are directly hooked up to this interface.

The IEC bus was designed for the serial shift register in the VIA (Versatile Interface Adapter) of the VIC-20 and its disk drive, but it turned out that the VIA had a bug that made the shift register unusable, so the IEC protocol had to be implemented in software. While the C64 has CIAs, in which the bug has been fixed, the 1541 still used VIAs. It wasn't until the C128 (in its native mode only) that the computer could talk to the floppy drive (Commodore 1570/1571/1581) in its intended speed.

1541 Disk Drive

The Commodore 1541 Disk Drive is the most common disk drive used with the C64. It uses 5.25" SS/DD (single side, double density) disks, but disks can be flipped, and the other side can be used as well, if the disks are double sided. The 1541 does not use the index hole, and uses software markers (SYNC) instead to be able to tell the start of a sector. Due to reliability problems of early drives, the 1541 only uses 35 out of the 40 possible tracks on a 5.25" disk. The tracks have a variable number of 256 byte sectors, ranging from 21 on the outside to 17 on the inside. The data is written in 4 different speeds. This makes an overall 683 sectors, or 174,848 bytes.

The file system is stored on track 18. Track 18, sector 0 contains the disk name and the BAM (block availability map), which stores one bit per sector (1 = free). Track 18, sector 1 is the first sector containing directory entries: There are eight 32 byte entries per sector, with a maximum filename length of 16 characters. The first two bytes of a directory sector point to the next directory entry sector.

The files on disk are also stored as a linked list. The first two bytes of every sector are either the track and sector number of the next block, or the first byte is 0 and the second byte is the number of valid bytes in this sector.

The 1541 is a stripped down version of the PET drive series, which had a parallel connection, and contained two 6502 CPUs: One for doing the filesystem and communicating with the computer, and one for reading data from disk and writing data to it, as well as encoding and decoding the data. The 1541 only has a single 6502 CPU running at 1 MHz, which (using timer IRQs) regularly switches itself between the two modes. The two virtual CPUs still communicate with each other using a messaging interface in the zero page. The 1541 has 2 KB of RAM at $0000-$07FF.

The 1541 has two VIA I/O controllers at $1800 (for the IEC bus) and at $1C00 (for the drive). The firmware is located at $C000-$FFFF.

Since loading an application or a game takes minutes on an unmodified C64, several ”floppy speeders” appeared (either as software on disk or built into applications, as ROM extension cartridges, or as internal replacement ROMs), that consisted of implementations of more optimized protocols for the IEC bus for both the C64 and the 1541. The 1541 code was uploaded using the old bus protocol. Such a new protocol would for example not do a handshake on every bit using the clock line, but shift a complete byte through in 4 steps, two bits at a time, using the clock and data line at the same time. This would of course only work if both CPUs were not interrupted. VIC timing on the C64 side could already affect this, so many floppy speeders turned off the screen while loading.

Other Peripherals

The 1541 is the necessary companion to a C64. It can be replaced by a 1570 (1541 with fast bus routines for the C128) or a 1571 (double-sided 1570), since they include a 1541 compatibility mode. The 3.5" Commodore 1581, which supports 880 KB per disk, can hardly be a replacement for a 1541, because most applications contain their own floppy speeders that make lots of assumptions on the exact on-disk format. For GEOS, it can be very useful though.

Creative Micro Devices sold and continues to sell a line of hard drives that have an IEC connection but contain a 3.5" SCSI drive inside. Although they have a 6502 CPU built in and allow code execution on their CPU, they are not compatible enough to replace a 1541 either.

Several memory extension cartridges exist for the expansion port of the C64: The Commodore REU (contains a DMA chip that transfers data between itself and main RAM), as well as the third party GeoRAM (maps a block to the $DE00-$DFFF area) and RAMLink (battery backed, designed as RAM disk).

Freezer cartridges like the Action Replay and the Final Cartridge were not only popular because they could dump all of memory to disk and thus copy certain copy-protected games, but also because they featured floppy speeders that disabled the original routines directly at startup time, without any effort from the user.

In the mid 90s, CPU speeders for the expansion port became popular. The 8 MHz Flash 8 is rare today, but many enthusiasts have a SuperCPU, which replaces the onboard CPU with a 20 MHz 65C816, which has a native 16 bit mode that can address up to 16 MB or RAM. There are a few applications and games that require a SuperCPU. The speedup of GEOS with a SuperCPU is significant.

In the 2000s, the enthusiast scene created all kinds of peripherals, like ethernet interfaces, IDE interfaces and SD card readers.

And there are not only peripherals: In 2004, the Commodore 64 DTV, a reimplemented C64 appeared in the form of a Joystick. The device is fairly compatible and can be extended to connect to a keyboard and a 1541.

Making

Manuel Odendahl

Algorithmic Music in a Box

Introduction

Since its conception in the early eighties, the MIDI bus has taken the world of electronic music by storm, and despite its limitations still remains an essential part of today's studios. The main attraction of the MIDI protocol for creative manipulation is that it represents symbolic information (notes, events, parameter manipulation, tempo information), rather than encoded audio. This allows us to use lightweight algorithms which can easily be implemented on small processors, and thus be used in DIY hardware. This paper does not go into specific implementation details, but provides an overview of MIDI DIY hacking.

There is a myriad of ways to use MIDI in a creative hacking context. MIDI often provides a way to poke into the internals of hardware synthesizer, so that their functionality, which is often cast in stone, can be extended by a clever outboard device. It also allows to create controllers that are customized for a certain kind of hardware synthesizer or performance approach, providing new ways of performing music and interacting with devices. Finally, custom music writing and generation tools can be built, allowing us to create algorithmically generated music that is tightly coupled to a custom hardware interface.

The purpose of this paper is to show in what ways a simple programmable MIDI controller with just a few basic controls (for example knobs, buttons or a joystick) can be used to completely change the workflow we have with existing MIDI hardware.

The MIDI Protocol

The MIDI protocol is a simple serial protocol implemented on top of a serial current-loop hardware bus. Data is sent at 31250 bps over a current loop, using a start bit and a stop bit for signaling, and 8 data bits. An optocoupler has to be used to receive MIDI information, while sending is done by simply toggling an output pin, sourcing 5 mA. The current loop is there to avoid ground loops when connecting multiple audio devices using MIDI. This means that interfacing standard microcontrollers to MIDI is very easy, as their UART component (most microcontrollers have one) can be used.

The MIDI protocol itself is very simple, and yet pretty cleverly built. The MSB of each byte is used to signal if the byte is a data byte or a status byte. Status bytes indicate the start of a command, which can be NOTE ON (0x90), NOTE OFF (0x90), CONTROLLER CHANGE (0xB0) and some more (the status bytes listed are those of interest to us in this article). The lower nibble of the status byte indicates on which "channel" the command is sent. Thus, each MIDI connection provides 16 distinct MIDI channels, which can be used to control different synthesizers for example.

Each status byte is usually followed by a few more data bytes (no MSB set) providing the actual information. For example, NOTE ON and NOTE OFF are followed by the note number and the velocity of the note. CONTROLLER CHANGE (which is used for example to send knob tweaking information) is followed by the controller number and the controller value. Most MIDI parameters thus have a resolution that goes from 0 to 127. There are quite a few extensions that allow for finer-grained parameters to be sent (NRPN), that allow relative parameters to be sent (relative CCs), as well as reduce the amount of data sent (running status), but in order not to burden the paper with technical detail, we will focus on just NOTE messages and CONTROLLER messages.

MIDI also has another class of messages called SYSTEM messages, which relate to all channels. These messages are comprised of SYSTEM REALTIME messages, which deliver clocking information, and SYSTEM EXCLUSIVE messages, which are pretty much a generic wrapper around binary data. MIDI CLOCK messages are a subclass of SYSTEM REALTIME messages. A clock event is sent at every 96th of the current tempo, allowing MIDI devices to synchronize to the tempo master. Additional SONG TRANSPORT messages allow the master to start/stop/position slaves inside a song structure. We will focus on tempo information in the "Tempo, Looping and MIDI Synchronization" part of this paper.

SYSTEM EXCLUSIVE messages are used by various manufacturers to control different manufacturer-specific of their devices. For example accessing more complicated parameters that can not be controlled directly using CONTROLLER CHANGE messages, storing and loading patches, polling the current user interface status and much more. We will explore SYSEX MESSAGES and what is possible with them in the "Hacking the Synthesizer User Interface" part of this paper.

Hacking the Synthesizer User Interface

This chapter shows a few ways in which the normal user interface of a synthesizer or hardware MIDI device can be extended through an external device, often completely changing the way in which the synthesizer can be used, often turning it into a completely different instrument.

Usually, SYSEX messages is the easiest way of accessing the internals of a synthesizer. One can often modify internal sound generation parameters that are not accessible through CONTROLLER CHANGE messages. A good example of this is the Elektron Machinedrum, which is the hardware synthesizer we have been focusing on. The MachineDrum provides 4 internal effects: Delay, Reverb, Equalization and Compression. However, the parameters for these effects can only be accessed using an internal menu, which takes about 5 key hits to access, and blocks the access to other parameters when active. Thus, tweaking the internal effects while performing is quite tricky and dangerous, and thus never done. However, the parameters can be accessed with SYSEX messages. We built a small controller that allows the user to directly control the parameters of the internal effects, which led to a completely new way of using the machine. With the controller, we were able to control the Delay Feedback parameter, allowing it to go into the dangerous feedback region knowing we were able to tame it at all times. This completely changed the approach we had to the machine, and paved the way for all the MIDI work we did later on.

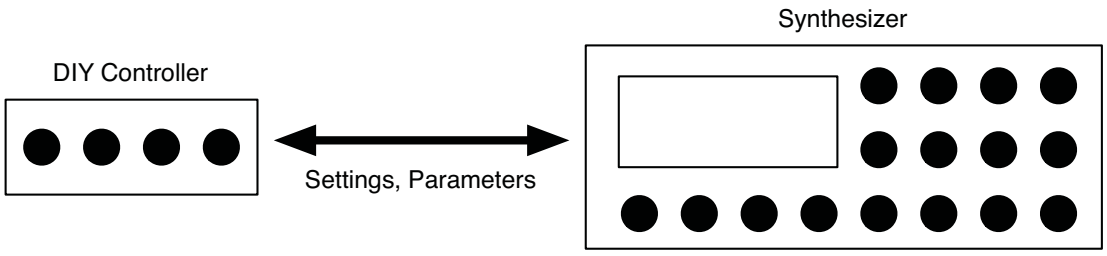

Another way we use SYSEX messages with the MachineDrum and MonoMachine synthesizers is loading the saved sound parameters (called a Kit), and analyzing them. The MachineDrum is a drum machine providing a plethora of drum and percussion sounds (each one of these called a Machine). Some of these sounds are pitched, but not in a conventional musical way (not following notes). Accessing

the SYSEX information, parsing the Kit settings, we were able to map conventional MIDI notes to the different pitch settings of the active Machines. This allows us to play a polyphonic snare drum using the keyboard, again completely changing the approach we have to the machine, as it is suddenly not a drum machine anymore, but a powerful and weird polyphonic synthesizer.

Being able to completely program our own controller allows for a host of tricks. For example, specific key combinations can be used to vary multiple parameters simultaneously (see also "Morphing Parameters and Exploring Parameter Spaces"), often driving sounds into very weird sonic regions. Accessing the saved settings over SYSEX, we are able to "restore" the original settings by just pressing a single button (something which would often require a few menu items on the synthesizer itself), allowing us to be very bold in our sound modifications.

Of course, much more prosaic things are possible through SYSEX access: archiving settings and patterns, merging patterns, storing sounds with useful names to be loaded later on (which is difficult to do on most hardware synthesizers), morphing between stored settings. A very nice feature is for example periodically storing the current pattern and current settings of a synthesizer while working on a song, and then later on allowing for a "Time Machine" kind of access. Undo memory is usually very limited or inexistant on a hardware synthesizer, sometimes making work very painful.

Morphing Parameters and Exploring Parameter Spaces

The number of parameters on synthesizers are usually pretty big, and even with a restricted number of parameters, the synthesis algorithms are pretty complex, allowing for a big range of different sounds to be created with a few basic ingredients. When adding complex modulation possibilities like LFOs, cross-synchronization of oscillators, multiple envelopes, the sound design space becomes virtually illimited. A lot of work has been done to come up with new ways of exploring the sonic spaces, for example the work of Palle Dahlstedt (http://www.ituniv.se/~palle/pmwiki/pmwiki.php).

A simple way to implement this kind of workflow on our devices is to make the controller do randomization of synthesizer parameters. For example, turning a knob on the controller can change a number of parameters inside a specific range, allowing the user to explore different combinations quickly without having to tweak every parameter by himself. This kind of modification becomes even more fun by using a different kind of hardware controller such as a joystick, or an acceleration sensor.

Modifying multiple parameters at once in a randomized fashion can however quickly lead to sonic chaos. A good way to tame the chaos is to use genetic mutation. A sound seed is planted and then evolved into different directions using the hardware controller. Once an interesting combination has been found, this can be made into a new seed and evolved again, or merged with another number of stored seeds. The VST Synthplant by Sonic Charge (http://www.soniccharge.com) or the genetic mutator of the Clavia G2 Modular Synthesizer work that way. Adding this kind of randomization to existing hardware synthesizers, maybe coupled to clever parameter mapping as described in the first part of this paper, breathes new life into old machines.

Tempo, Looping and MIDI Synchronization

Most music, and especially electronic music has a firm rhythmic foundation. This is also visible in the MIDI specification, providing SONG START, STOP and MIDI CLOCK messages. The MIDI CLOCK message (0xF8) is periodically sent on every 96th note, and has precedence over every other message,

thus allowing clock synchronization to be tight. The receiving side can either clock its internal clock directly from the MIDI CLOCK messages, or use a phase-locked loop to generate its own clock. Once two devices are synchronized, they run at the same tempo, and the slave devices follows the tempo of the master device.

Implementing tempo synchronization in our devices allows us to take the user interface modification to another level, as we can now synchronize our modifications to the current tempo. This allows us for example to loop parameter modification for a rhythmically sound duration (for example 4 bars). It also allows us to cue parameter modifications to become active on a specific beat.

Having a synchronized clock also allows us to generate LFOs for specific parameters. An LFO is a low-frequency oscillator that slowly changes a parameter according to different waveforms. For example, it can make a parameter go up and down every 2 seconds over a range of 40 units. A clock synchronized LFO allows us to vary sounds in a rhythmically interesting manner (for example allowing for polyrhythms on a sonic level). The standard way to provide this kind of parameter manipulation is to use the automation features of a sequencing program on the computer, and either record knob tweakings or draw them into a timeline. This is quite time-intensive and doesn't allow for the kind of exploration that is easily done on hardware.

Often, hardware synthesizer only provide a limited number of LFOs, or with limited waveform possibilities. Adding our own external LFOs allows us to extend the possibilities of the hardware synthesizer while keeping the amount of work involved low. It also allows us to explore the possibilities of an LFO in itself, for example: replacing the waveform with a table of values, synchronizing two LFOs together, locking the LFO values to interesting values (for example tonality dependent, or rhythmically sound).

Another example combining the tempo synchronization and the parameter hacking is a hack for the Elektron MachineDrum. The MachineDrum allows to load and record samples, and move their start- and endpoints parameters (where the sample starts and where it ends). Assuming we have loaded a rythmic sample (for example a drumloop), moving the playback points in units of 8 allows us to start the sample at different rythmically sound values, thus "chopping" the sample, also allowing for interesting reverse sample effects and pitchdown/pitchup effects. This kind of manipulation is the basis for a lot of electronic music, for example drum'n'bass, breakcore, IDM. However, the standard interface of the MachineDrum makes it impossible to tweak the playback points in this way, and also to modify the parameters on rythmic hits. Using the external controller, we are able to move the endpoints and startpoints in rhythmically sound values (units of 8 for 16th notes, 16 for 8th notes, etc...), and modify the parameters only on rhythmic events and not in between. This allows for tight live drumloop chopping. We can take the concept further and implement automatic drumloop-chopping algorithms such as those invented by Nick Collins (http://www.cogs.susx.ac.uk/users/nc81/).

Algorithmic Rhythms

Taking tempo synchronization further, we can actually implement complete sequencers on our DIY hardware. A simple kind of sequencer that lends itself well to a hardware interface is an arpeggiator. An arpeggiator takes a number of input notes (a chord), and plays back notes from this chord synchronized to a clock, thus turning the chord into a quick arpeggio. Mapping parameters of this arpeggiation to hardware controls allows to create interesting rhythmic and harmonic texture, playing with

note length, note range, and different other parameters. For example, something "common" arpeggiators often don't provide is to couple synthesis parameters to the arpeggiation. This can be easily added to our external controller, allowing us to play with accents, with rhythmic modification of synthesis parameters, once again turning our external controller into a completely new instrument of its own.

We can also devise new sequencer ideas. A very fascinating algorithm to generate rhythms is the euclidean algorithm, which can be used to generate most of the western and african rhythms (see the paper The Euclidean Algorithm Generates Traditional Musical Rhythms by Godfried Toussaint). Using just 3 numbers (number of hits, length of pattern and start offset), we can generate complex and natural sounding polyrhythms. Implementing this algorithm on our hardware controller, we can very quickly and easily modify and explore different rhythms, without having to actually calculate the rhythms and input them into a normal sequencer.

Of course, this can be extended to any musical theory that lends itself well to an algorithmic computer-based implementation. Having direct haptic access (through the controller interface) to the sequencing parameters allows for a much more indepth and natural exploration of different kind of algorithmic musical processes.

Conclusion

We hope that we were able to show the immense range of possibilities that a DIY MIDI controller offers in terms of musical exploration. It allows us to extend (hack) existing MIDI devices, adding custom functionality. It allows us to explore musical possibilities through an intelligent and customized interface. It extends the performance possibilities by allowing rich tempo-synchronization possibilities, and finally permits us to implement, prototype and interact with complex algorithmic sequencing models. Videos of these features as well as schematics and sourcecode can be found on our website at http://ruinwesen.com/ .

Solar-powering your Geek Gear

Michael "script" Arndt

December 27th 2008

This article will show you how to solar-power your laptop, PDA, cell phone, portable fridge or almost any other small device. It explains how to choose the right solar panel, how to use (or not use) a voltage regulator and why it might be useful to buffer the energy. It introduces a small and quite simple device to measure power and energy savings and finally some typical applications are discussed.

1 Introduction

1.1 Warning

This project involves currents and voltages that might be considered dangerous. Everything that is discussed here should be used at one's own risk, connecting devices to anything else than their dedicated chargers could easily break them, void their warranty or cause even worse to happen. There is absolutely no warranty from the author, you have been warned.

1.2 Motivation

Why would someone want to power his devices using a solar-panel? There are a couple of good reasons to do so. First of all, solar power is just a fascinating thing - you put the panel into the sun, plug it in and (hopefully) it is just working - no battery, no power plant necessary. Being independent of power sockets, there are a lot of scenarios where this technology could be used: during camping in the wilderness, lonely on the top of some mountain, in the park or in the garden to name just a few places.

You might also think that it might be possible to save a lot of money using your solar panel. Unfortunately I have to disappoint you for now. At the time of writing this article, you would have to use it *really* long to save money, but as prices for photovoltaics might drop and energy-prices might rise in the next few years this could of course change.

Talking about money, let us consider how much you have to spend. Currently you can buy solar-panels for about 50 EUR (about 60 USD) per 10 Watts. The panel that has been used during research for this project was about 200 EUR for 40 Watts. Depending on what you want to power (see Section 2.2) you will also have to spend money for e.g. voltage regulators, capacitors or plugs.

2 The Solar Panel

2.1 Characteristics

Solar-panels can be characterized by their short-circuit-current I_{sc}, their open-circuit-

voltage V_{oc} and their maximum-power-point (MPP) (V_{mpp}, I_{mpp}). While V_{oc} is almost constant between several lighting-conditions the current I_{sc} is heavily dependent on radiation and cell-temperature (see also [4]).

Figure 1 shows two voltage-current-curves under different lighting-conditions. As illumination increases (decreases) the short-circuit-current I_{sc} and thus the available maximum power P_{mpp} also increases (decreases).

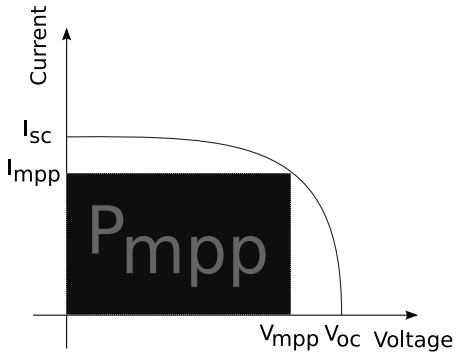

(a) Testing-conditions (1000W/m^2, 25°C)

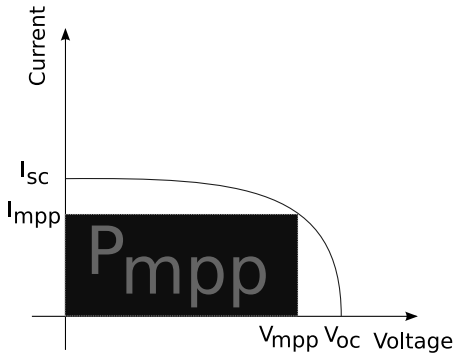

(b) More typical conditions

Figure 1: Solar-panel under different lighting-conditions

Looking at the figure, it is obvious that you need to include some safety-margin when choosing your panel.

2.2 Deciding what to power

As mentioned before, there are a lot of small devices that can be solar-powered. Some devices are very picky about their input-voltage while others are not. The first thing to do is to have a look at the ratings of the device or its original power-supply. The continuous power of the device should not exceed the power of your panel including some safety-margin (see above). Higher peak power-ratings can possibly be buffered using a capacitor (see section 3.3).

Lower voltages than stated in the rating are often okay while higher voltages should definitely be avoided if you do not want to break your device (or know it very well and you are sure that it will be alright). Devices that are rated below the open-circuit-voltage of the panel should be powered using an adjustable voltage regulator. The regulator will take care of limiting the voltage so that the device will not be damaged. In some cases it may also be necessary to limit the current through the device, see section 4.2 about Nokia phones for a prominent example. Despite of that, the majority of devices will do fine with just a voltage-regulator, see section 3.2 for an efficient wide-range voltage-regulator.

An adjustable laboratory power-supply is a valuable tool to test your devices using different voltages and to measure power consumption. If the device is not working properly under these testing-conditions it is very likely that it will not work using the solar-panel either.

3 Accessories

3.1 Connectors

To give a robust and safe plug-and-play experience special connectors have been cho-

sen for the project. The following requirements had to be satisfied:

1. Connectors should be common for all devices
2. No shorting and no reversal should be possible
3. Connectors should support a current of up to 10A

After some research I came up with the RIA Connect 230/249 Series [6] that are distributed in Germany e.g. by Reichelt[1] [5]. See figure 2 on how these connectors look like.

Figure 2: Connector cord

Having chosen the connectors some useful adapters and tools can be build. Older ThinkPad-Adapters can be build using simple DC plugs[2] with the following specification: $\emptyset_i$ = 2.5mm (inner diameter), $\emptyset_o$ = 5.5mm (outer diameter). Newer ThinkPads (20V) have a more complex plug which is harder to obtain, one solution would be to cut it from a cheap replacement power supply. The same applies for cellphone-connectors. 12V Car-Equipment, which includes several chargers and inverters, can be connected using the appropriate sockets. To connect several plugs together, a small distribution board can be built (see figure 3).

Figure 3: Distribution board

3.2 Universal Voltage Regulator

Typical voltage regulators only work with input voltages higher than output voltages (step-down), but there are also step-up-converters that convert lower to higher voltages. For this project, a wide range of output-voltages is required and the panel provides a wide range of input-voltages under load. This is the reason why I have chosen a very special voltage-converter - it is based on the LTC3780 (see [7]), a buck-boost-controller which allows seamless switching between Step-Down (Buck) and Step-Up (Boost) modes. The regulator can be bought as "USW-525" at [3], see Figure 4. It has the following specs and can thus be used for a board range of applications:

- Input-Voltage: 7-25V
- Output-Voltage: 4-25V
- Output-Current: up to 5A
- Efficiency: up to 97%

[1] Reichelt Part-No: AKL 249-02, AKL 230-02
[2] Reichelt Part-No: HS 25-9

Using this regulator, I was able to power older ThinkPads (16V, 4.5A), 12V-car-equipment like chargers and an inverter as well as many other devices.

Figure 4: ELV "USW-525" universal voltage regulator in a case

3.3 Buffering the Energy

Many devices have short peaks of very high energy consumption, see figure 5 for an oscillogram of peak power consumption with a ThinkPad X300 while switching desktops in Fluxbox.

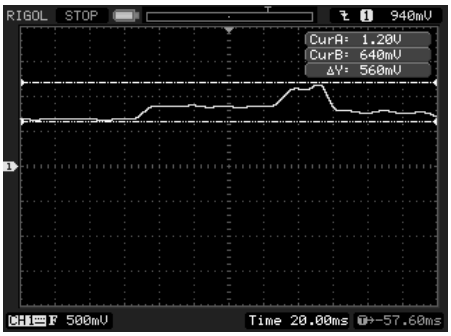

Figure 5: Peak power consumption, 1V $\widehat{=}$ 1A across a 1Ω-Shunt, $\Delta I = 0.5\text{A} \Rightarrow \Delta P = 10\text{W}$

In case of a laptop this peak-consumption can be buffered using the internal battery while other devices that do not contain a battery might just reset or power off. The solution to this problem is to use a *very* big capacitor to act as a buffer to the solar-panel. Depending on the capacity and the load it might also be able to buffer a temporary shading of the panel (e.g. when someone is walking in front of it).

Figure 6 shows a 1 Farad power capacitor that is usually used for car-HIFI. Please note, that this capacitor can generate *extremely* high currents! Don't ever short it when charged and never try to charge it directly using a high-current source (e.g. a battery)! You should really put a fuse into the cable and isolate the terminals.

Figure 6: Fused powercap

Additional care must be taken that no current can flow backwards from the capacitor to the solar-panel (the same that applies when charging batteries directly with solar-panels). This can be achieved by using a blocking diode between the panel and the capacitor.

3.4 Measuring Power and Energy

Soon after some simple test-setup with the solar-panel and a laptop worked, I wanted to know what power the panel is currently delivering and how much energy I could save in one afternoon.

For this purpose, I built a small device consisting of an Atmel ATmega8 micro-controller, a display, an operational amplifier as well as some other components.

The input voltage is measured directly using a voltage divider and the current is measured using a small shunt of about 10mΩ. A current of 2A will thus cause a voltage-drop of about 20mV at the shunt. This voltage is amplified by the operational amplifier and fed into the analog-digital-converter of the micro-controller. This drop is acceptable - on the other side, if no amplifier had been used but instead a 1Ω-shunt, a current of 2A would cause a voltage-drop of two precious volts.

The software reads out the analog values, converts them to voltage and current values and calculates power ($P = V \cdot I$) as well as energy ($E = \int P \mathrm{d}t$). All values are shown on the display. If input voltage falls below a threshold of about 9V, the device enters shutdown-mode, which means it will write the energy-value to EEPROM and disable further measurements (until it is wakened by the button). This strategy helps to prevent EEPROM corruption when power fails as well as exaggerated EEPROM-writes.

A more detailed description of the device as well as schematics and software can be found at my personal website [1].

4 Applications

4.1 Laptop

Laptops usually have a built-in battery which makes them quite easy to use with solar power. As soon as voltage drops due to higher load or cloudy sky, the laptop can switch to its internal battery and you can continue working. Power consumption of modern laptops is around 10-30 Watts, but while charging its battery and can be *a lot* more (e.g. 60W), so what you need to do is to prevent the machine from charging (as long as you do not have a very large panel that can delivery that additional power).

ThinkPad Laptops can be instructed to not charge the battery by using the `tp_smapi`-interface. The following will set the battery-start-threshold to 10%, so that the battery will not be charged as long as it is above 10%:

```
echo 10 > /sys/devices/platform/\
   smapi/BAT0/start_charge_thresh
```

Apple's iBooks and Powerbooks have a Sense-Pin in their connectors that can be used to prevent battery-charging, please see the technical document at Apple's site for more information [2].

Unfortunately I was unable to find more information about other laptops, so you will need to figure out a way to prevent it from charging the battery. Physically removing the battery is no good option, unless you have a good energy-buffer (see section 3.3) and know that there will be no clouds soon.

One strategy is to apply a voltage that is low enough, so that no "AC connected" will be detected (and thus no charging will take place) but still high enough to entirely power the laptop. Results may of course vary and depend on the implementation of

your vendor - I was at least able to verify this procedure with the already mentioned ThinkPad X300.

4.2 Phone

Nokia phones will not charge with a constant voltage if the current that the source can deliver is too high. Reasons why this happens to be so are unknown, but probably the internal charging-circuit has no means of limiting the current flowing to the battery.

Original chargers have their current limited to $\approx$ 800 mA, so if you want to charge a Nokia phone you will have to limit the current. The phone will test the available current at the begin of the charging-procedure, if it is too high it refuses to charge. The most simple current-limiter-circuit is just a resistor in series to the constant voltage source, see Figure 7. The voltage of the source (e.g. your voltage-regulator) should be set to about 5-6V (see the back of the charger), so according to Ohm's law, a resistor of $R = \frac{U}{I} = \frac{5V}{800mA} \approx 6\Omega$ will do the trick. Experiments showed that 5Ω also worked with several phones.

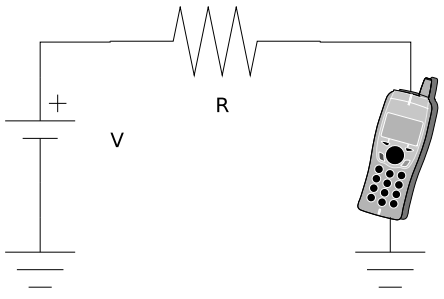

Figure 7: Current-limited charging

4.3 Electrical Fridge

Electrical fridges are quite simple devices, most are rated around 40 Watts (at 12 Volts). You could of course set your voltage regulator to 12V and hope that your solar-panel is strong enough, but those fridges will most likely also work with voltages below 12V (they will not get as cool, though). So if you use a panel that provides about 40 Watts or less you can just try plugging them directly together, the panel will then act as a constant current-source and limit the voltage at the fridge.

5 Conclusions

A rather small solar-panel (40 Watts), combined with some accessories like connectors, adapters, a voltage regulator and a capacitor can be used to have a lot of fun with a variety of small devices. Powering laptops that allow battery management, works great and enables the user to work for hours without the need for power sockets or heavy batteries. Further research should be done on laptops that cannot easily be prevented from charging, however, as charging batteries uses a lot more power.

References

[1] Michael Arndt, "Personal Website", http://scriptkiller.de/, December 10 2008.

[2] Apple Computer, "Apple Technical Q&A QA1266, Power adapter Sensing for the 17" PowerBook", http://developer.apple.com/qa/qa2001/qa1266.html, December 10 2008.

[3] ELV, "ELV Universal-Step-up / Step-

down-Spannungswandler USW 525", http://www.elv.de/, December 10 2008.

[4] Vijay Raghunathan, A. Kansal, J. Hsu, J. Friedman, Mani Srivastava, "Design considerations for solar energy harvesting wireless embedded systems", in *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on.* April 15 2005, pp. 457 – 462.

[5] Reichelt, "Reichelt Online-Shop", http://www.reichelt.de/, December 10 2008.

[6] RIAconnect, "RIAConnect Website", http://www.riaconnect.com/, December 10 2008.

[7] Linear Technologies, "LTC3780 Datasheet", http://www.linear.com/, December 10 2008.

Science

about cyborgs and gargoyles - state of the art in wearable computing

Kai Kunze
Embedded Systems Lab, University Passau
http://wearable-computing.org/
http://wearcomp.eu/
http://twitter.com/kai_ser/
{first.lastname@uni-passau.de}

This paper gives you an introduction to wearable computing. It starts of with a small introduction concerning the terms, then moving over to some concrete application scenarios, from a maintenance to firefighters. Subsequently, we discuss some tools used to build wearable, context-aware systems. Finally, we take a look at some problems concerning the current systems and privacy implications.

Track [**Science**]: Language—*English*

1. INTRODUCTION

With the widespread integration of embedded processors and the proliferation of mobile devices, the vision of pervasive/ubiquitous/wearable computing becomes more and more feasible: computing that helps us transparently in our everyday life, focusing on the our needs without disturbing us. In the following pages, I try to give a brief overview of the field and then go into details about some of our efforts towards this vision using concrete application examples.

Personally, I do not see too much differences between the visions outlined in ubiquitous, pervasive, wearable computing and ambient intelligence. Each discipline just focuses more or less on particular aspects of the technology.

Wearable computing is a nice term as it describes the ideas in a simple way. Computing as useful and unobtrusive as clothes.

The pervasive/wearable computing field is as such interdisciplinary. It combines aspects of embedded systems, signal processing, artificial intelligence, machine learning and human computer interaction.

Subsequently, I will discuss some of the basic assumptions of the field, giving some application scenarios we work on at the Embedded Systems lab. Afterwards, I give details on how current pervasive systems are prototyped and implemented. Finally, I go into some of the technical challenges we still face in this field.

Of course, all of the technologies introduced in the paper have some pretty strong privacy implications. On the one hand, computing that knows what the user is doing can be a bless, as it is able to help him to accomplish the tasks he is working on. On the other hand, it can track and record every move the user does and therefore implement an Orwellian Big Brother. I will shortly rough out some of the ethical issues in the last section. The content of this paper is a summary of papers given in the bibliography.

Regarding the title 'About cyborgs and gargoyles', sadly, I have not much background in implantable sensors, therefore I won't discuss them here. Although implantables are extremely interesting, so far, the user acceptance is very low except in the medical domain. The title is, I think, still justified as having wearable technology to support us is the first step to integrate computing seamlessly in our life, as extensions to our sensory inputs.

2. CONTEXT RECOGNITION IS THE KEY

We basically want use the already existing substrate of microprocessors and sensors embedded in our environment to infer relevant information about the users current activities and state. This "relevant information" is often referred to as context. Detecting this relevant information is then defined as 'context recognition'. There are some debates in the research community on what context actually is. For the remainder of this paper, I will use the very practical, broad definition given above. We have the infrastructure in place given from today's embedded systems that helps us gather information about our current activities and state. There are plenty of papers([5,9,12]) describing context recognition systems. The sensors used for such systems vary substantially, from video cameras used for computer vision, over microphones to inertial motion sensors. One can classify them tough into two distinct types:

—infrastructure sensing: Stationary sensors like a microphone array, a stationary camera, a radio frequency localization system like Sputnik etc.

—onbody sensing: sensors worn on the user's body, for example: inertial motion, (accelerometers, gyroscopes etc.), galvanic skin response sensors etc.

Using a combination of these sensors as information sources, we should be able to support any high level user acitvity, from work and collaboration over everyday living to healthcare and game/fitness applications.

2.1 application scenarios

The idea to use context recognition systems, that help the user in everyday activities seems promising. Yet, how do such system examples look like? In this section, I describe some concrete application scenarios from research projects implemented by the Wearable Computing Lab, ETH Zuerich [1] and the Embedded Systems Lab [2], Uni Passau. Most of them are taken from the WearIT@Work project, an EU Project with over 50 partners in academia and industry (see [10] and http://www.wearitatwork.com/) and RELATE, a smaller EU project focusing on baseline research (see http://eis.comp.lancs.ac.uk/relate/).

2.2 search and rescue

The first application scenario we look at is very challenging, supporting firefighters during rescue operations. Today, firefighters use ropes in cases of impaired visibility to mark paths taken into buildings to help them in finding the exit again. They call these ropes 'lifelines'. For many situations, this technique works reliably well. Yet,

[1] http://www.wearable.ethz.ch/
[2] http://esl.fim.uni-passau.de

Figure 1: A first generation RELATE sensor node placement on the floor of the training facility and a firefighter equipped with several sensor nodes

there are also serveral limitations to the rope 'lifeline'. The lifeline can get stuck or be cut under doors or other objects and it limits the effective range of the firefighters. In the WearIT and RELATE projects, we developed the use sensor networks for a virtual lifeline. Such that, firefighters automatically deploy sensor nodes along their paths, using a dispenser mechanism, establishing an ad-hoc infrastructure for positioning, sensing and communication. The firefighters can interact with this sensor network by way of wearable computing equipment and receive navigational information on e.g. a head-mounted display, over a headset or tactile feedback. The technology is to enable automatic tracking of a building's floor plan, tagging dangerous areas, checking rooms that have been searched for unconscious persons etc. Although we showed a first prototype implementation at Ubicomp 2007 [6], we are still far away from a system that is able to support firefighters in a real emergency.

2.3 healthcare and hospitals

The next application area, we take look at is in the hosptial and pervasive healthcare area. As you know, hospitals are a highly complex environment. Although there is already a quite some IT infrastructure at place, traditional interaction with a computer system is not really possible during most medical procedures. One of the problem areas my colleagues tackled over a period of 3 years, is the daily ward rounds of doctor and nurse teams. Currently, most of the interaction is still on paper.

In the system my colleagues proposed, implemented and tested, the doctor and nurse teams move around without paper,laptop or filing card. The doctor wears a wrist band (inertial motion sensor + rfid reader) and a qbic belt intergrated linux pc, depicted in Figure 2. The doctor identifies the patient by means of his RFID reader on his wrist. The relevant patient files appear on the screen, attached to the bed. The doctor immediately gets an overview of all the important data about the patient, With the inertial sensor attached to his wrist, he can navigate through the application without a mouse or keyboard. This is keeping his hands sterile and

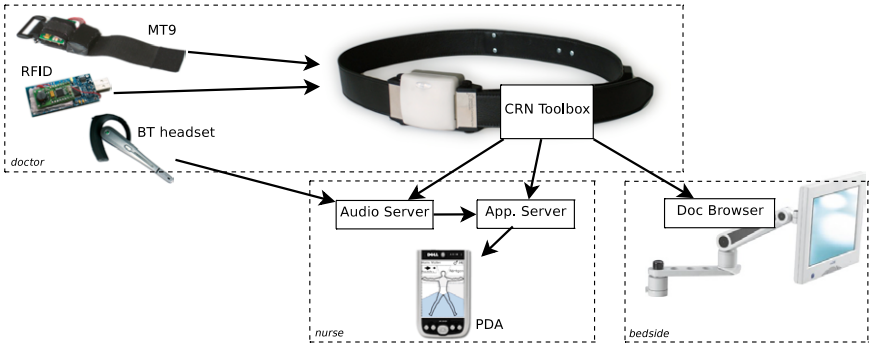

Figure 2: System architecture for ward rounds

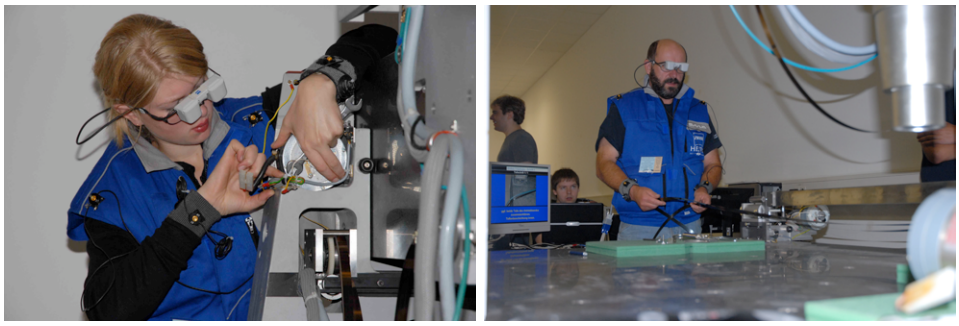

Figure 3: Pictures from the experimental sessions for the maintenance scenario. From left to right: Installing the xmotor and the ygears.

free to examine the patient at any time.

The basic system architecture is given in Figure 2. The system was evaluated over several weeks in a hospital in Steyr, Austria, using real doctor/nurse teams and patients (see [1]).

2.4 maintenance

Another application focus of our work is the domain of wearable maintenance systems.There are many such systems proposed and implemented since the early days of wearable computing. These systems aim to provide maintenance personnel with access to complex electronic information with as little interference as possible to the primary task at hand. Typically, they rely on head mounted displays (often with augmented reality), input modalities that minimize hand use (e.g. speech, special gloves) and interfaces that aim to reduce the cognitive load on the user. It is widely believed that wearable maintenance systems can benefit from automatic work progress tracking. Main uses for such tracking are just in time automatic delivery of information (see the manual page that you need without having to explicitly demand it), error detection (e.g. you forgot to fasten the last screw), and warnings (e.g. do not touch this surface).

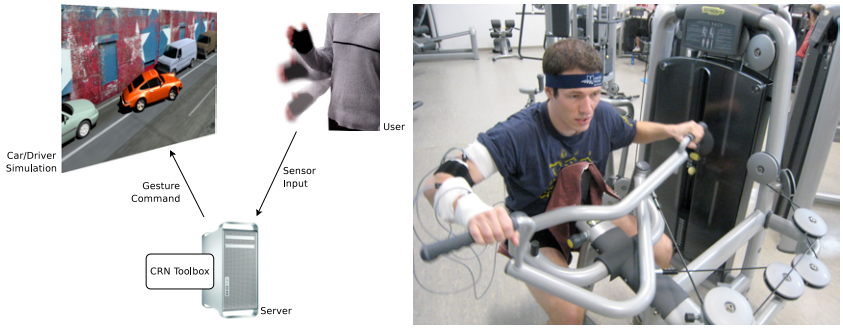

Figure 4: Car parking game, You are the person helping a virtual driver to park a car using hand gestures and a picture from data recordings to recognize gym exercises

In a paper accepted for pervasive 2009, we show a quantitative study on the usefullness of context recognition in a real world maintenance environment [8].

2.5 games and sports

Of course, wearable systems can also support gaming and sports applications. Although there are already several gaming platforms on the market that follow this apporach (Sony's eyetoy, Nintendo's Wii). Their 'recognition' capabilities are very limited and users cannot use their expertise gathered by doing the real sport in game. We have shown how to recognize difficult gestures online to enable car parking game [2]. Some of our efforts in this area focus on the recognition of martial arts movements and gym exercises.

3. HOW DOES IT WORK?

Today, the development of context recognition systems is mostly done in two phases. First the recognition method (sensor setup, feature set, classifiers, classifier parameters, fusion methods etc.) is designed. Researchers apply algorithms from the fields of signal processing, machine learning, artficial intelligence. In this first phase experimental data is mostly fed offline into conventional rapid prototyping tools such as Matlab. These tools provide a rich reservoir of 'off the shelve', parameterizable algorithms and visualization methods. Thus, different system variants can be tested quickly without the need for time consuming implementation work.

Unfortunately, most such simulation environments are not suitable for actually running applications, especially in mobile and pervasive environments. In general, they depend on custom 'engines' or libraries requiring large memory footprints and high computing power. Consequently, the implementation of activity recognition applications is mostly done in a separate, second phase. The selected algorithms are implemented in an appropriate programming language and then distributed to specific devices.

We actively develop the Context Recognition Network (CRN) Toolbox ([4] and [3]) (available under LGPL from http://crnt.sf.net) to combine the two phases and to permit quick construction of complex multi-modal context recognition systems, that can be immediately deployed in the targeted environment. The CRN

Toolbox has been used to develop the protoypes for the application scenarios discussed above.

3.1 practical tips to build context recognition systems

First of all, you can obtain the source code and examples of some of the systems presented above by checking out the toolbox source from http://crnt.sf.net . Here you can also find some more information about the supported platforms. You should get it working on all systems that have POSIX Threads. However, please understand that the CRN toolbox is in the very first stages of development, this means that documentation and examples are still lagging.

If you are in genereal interested in signal processing and machine learning algorithms, you might want to take a look at the tutorials of the 'Intelligent Systems' course I teach this semester at the Unversity Passau (http://esl.fim.uni-passau.de/wiki/index.php/Intelligent_Systems_Winter_2008).

Unfortunately, most of the tutorials still use Matlab, which is quite expensive and therefore not so easy to obtain for a hobby researcher. Yet, I started to include the equivalent scripts in python using scipy. We are in process of switching to it for the tutorials, although it might take a while until all of the scripts are ported. For people interested in free data exploration tools, I can recommend the following:

—Scientific Python, great package with nearly everything you need to get going http://scipy.org/

—Matplotlib, nice plotting library. http://matplotlib.sourceforge.net/

—Sage uses scipy, matplotlib and other packages good alternative to maple, matlab etc. http://www.sagemath.org/

The weka machine learning package in Java is also a good toolkit for alogrithms.

4. BECOMING OPPORTUNISTIC

Another issue away from rapid prototyping and easy testing, we are currently struggling with, is the augmentation problem. Most of the systems introduced above work well if you use dedicated sensors with a fixed placement and known orientation. Most inference algorithms cannot cope with failing or displaced sensors. This is actually my personal field of interest and the focus of my phd. work. There are a couple of interesting approaches to get over these limitations (for example [11]). In [7] we show how one can recognize the surface on which a mobile phone is placed using the virbration motor and frequency beeps. Unfortunately, these algorithms are not yet very practical, as they still run offline using scipy and matlab and cannot be implemented on commodity embedded hardware.

5. PRIVACY IMPLICATIONS

As this is a technical paper, I cannot go much into the ethical issues and problems associated with this technology. Also, this is not really my expertise. Yet, anybody will agree that wearable computing raises some severe ethical problems. On the one hand, the system that can help a user avoid mistakes and errors is able, on the other hand, to report error statistics to his employee etc. I want to raise awareness to these issues, as a lot of people are uneasy with the installment of video cameras

or similar. I believe that camera surveillance is not so much a danger to our privacy, than much simpler sensors and technology that can be easily embedded in mobile phones or other household devices. This is partly due to the fact, that computer vision has severe limitations (light, angle etc.) compared to the fusion of several, more simple "sensor" signals. For example, to track a person's purchasing habits in the city, it is way easier to use the credit card record combined with the cell log of the mobile phone than cameras distributed downtown.

REFERENCES

[1] K. Adamer, D. Bannach, and P. Lukowicz. Developing a wearable assistant for hospital ward rounds: An experience report. In *Proceedings of the International Conference for Industry and Academia on Internet of Things*, 2008. To appear.

[2] D. Bannach, O. Amft, K. Kunze, E. Heinz, G. Tröster, and P. Lukowicz. Waving real hand gestures recorded by wearable motion sensors to a virtual car and driver in a mixed-reality parking game. In *Proceedings of the IEEE Symposium on Computational Intelligence and Games*, pages 32–39, 2007.

[3] D. Bannach, O. Amft, and P. Lukowicz. Rapid prototyping of activity recognition applications. *IEEE Pervasive Computing (to appear)*, April-June, 2008. to appear.

[4] D. Bannach, K. Kunze, P. Lukowicz, and O. Amft. Distributed modular toolbox for multi-modal context recognition. In *Proceedings of the 19th International Conference on Architecture of Computing Systems*, volume 3894 of *LNCS*, pages 99–113. Springer, 2006.

[5] N. Kern, B. Schiele, H. Junker, P. Lukowicz, and G. Tröster. Wearable sensing to annotate meeting recordings. *Personal and Ubiquitous Computing*, 7(5):263–274, 2003.

[6] M. Klann, T. Riedel, H. Gellersen, C. Fischer, M. Oppenheim, P. Lukowicz, G. Pirkl, K. Kunze, M. Beuster, M. Beigl, O. Visser, and M. Gerling. Lifenet: an ad-hoc sensor network and wearable system to provide firefighters with navigation support. In *Adjunct Proc. Ubicomp 2007*, pages 124–127, September 2007.

[7] K. Kunze and P. Lukowicz. Symbolic object localization through active sampling of acceleration and sound signatures. In *Proc. 9th Int. Conf. on Ubiquitous Computing - Ubicomp 2007*, Innsbruck, Austria, 2007.

[8] K. Kunze, F. Wagner, E. Kartal, and P. Lukowicz. Does context matter ? - a quantitative evaluation in a real world maintenance scenario. In *Pervasive Proceedings*, 2009. To appear.

[9] J. Lester, T. Choudhury, and G. Borriello. A practical approach to recognizing physical activities. *Proceedings of Pervasive*, Jan 2006.

[10] P. Lukowicz, A. Timm-Giel, M. Lawo, and O. Herzog. Wearit@work: Toward real-world industrial wearable computing. *IEEE Pervasive Computing*, 6(4):8–13, Oct-Dec 2007.

[11] D. Minnen, C. Isbell, I. Essa, and T. Starner. Discovering multivariate motifs using subsequence density estimation and greedy mixture learning. *Artificial Intelligence Proceedings*, Jan 2007.

[12] C. Randell and H. Muller. Context awareness by analysing accelerometer data. *The Fourth International Symposium on Wearable Computers*, 1(1):175–176, 2000.

Climate Change - State of the Science

by Prof. Stefan Rahmstorf
Potsdam Institute for Climate Impact Research (www.pik-potsdam.de/~stefan)

Some basic facts about global warming

Important core findings of climate research have been so well confirmed in recent decades that they are now generally accepted as fact by climate researchers. These core findings include the following:

1. The atmospheric CO_2 concentration has risen strongly since about 1850, from 280 ppm (a value typical for warm periods during at least the past 700,000 years) to over 380 ppm.
2. This rise is entirely caused by humans and is primarily due to the burning of fossil fuels, with a smaller contribution due to deforestation.
3. CO_2 is a gas that affects climate by changing the earth's radiation budget: an increase in its concentration leads to a rise in near-surface temperature. This has been known since the 19th Century and is well-established physics. If the concentration doubles, the resulting global mean warming will very likely be between 2 and 4°C (the most probable value is ~3ºC), with the remaining uncertainty due to climatic feedback effects.
4. Since 1900, global climate warmed by ~0.8°C. Temperatures in the past ten years have been the highest since measured records started in the 19th century and for many centuries before that (Fig. 1).
5. Most of this warming is due to the rising concentration of CO_2 and other anthropogenic gases. These would in fact explain more warming than is observed, were they not offset in part by the cooling effect of aerosol pollution (smog).

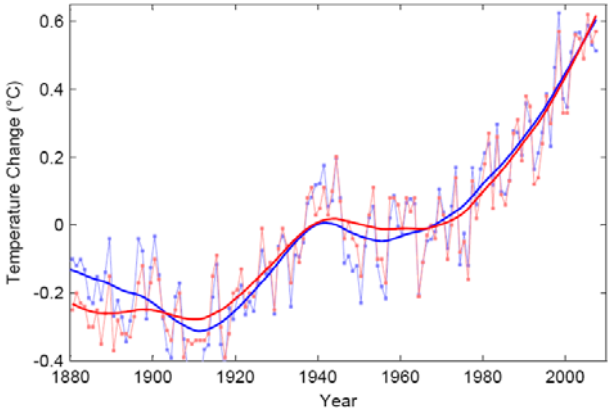

Fig. 1. Global mean temperature up to 2007 according to the two main data compilations: NASA in red and Hadley Centre in blue. Dots show annual values, heavy lines a non-linear trend smoothed over 15 years. Numbers are deviations relative to 1951-1980.

These findings are based on decades of research and thousands of studies. The extraordinary consensus reached is seen in the statements of many international and national professional bodies which have extensively and critically assessed the scientific evidence. In addition to the well-known reports of the IPCC, there are public statements of the National Scientific Academies of all G8 countries, the American Geophysical Union (AGU), the World Meteorological Organisation (WMO), the scientific Advisory Council on Global Change (WBGU) of the German government, and many others. These organisations have again and again come to the same key conclusions.

From points 1. – 3. follows that a further increase in CO_2 concentration must lead to a further rise in global mean temperature (Fig. 2). For a range of plausible assumptions about future emissions, by the year 2100 this rise will reach 2 - 7 ºC above preindustrial values.

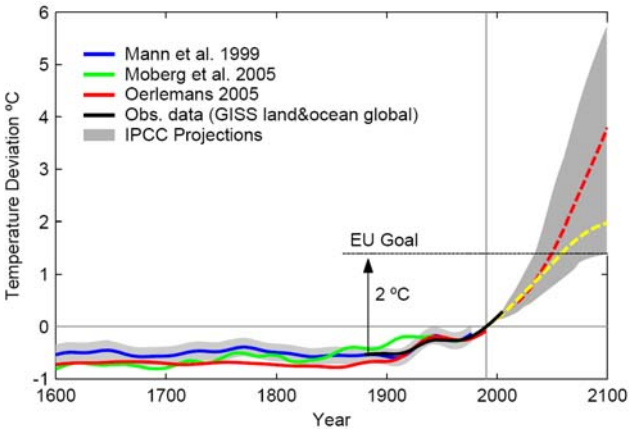

Fig. 2. IPCC projections for global mean temperature in the 21st Century. The grey band shows the full range of scenarios; red and yellow are two examples (B1 and A2). For comparison, several reconstructions for temperatures of the past centuries are included. The EU 2º-limit is also shown.

For comparison: the last major global warming was the end of the last great Ice Age (about 15,000 years ago); it involved a global warming of ~5ºC over a time span of 5,000 years. Unchecked anthropogenic warming could reach a similar magnitude over a fraction of this time – and, of course, starting from an already warm climate.

Impacts and risks

Whether this warming is considered a "dangerous" climate change can, of course, not be determined by scientists alone, as it depends on a societal value judgment about what is dangerous. However, science can help to clarify what are the risks that arise from such unprecedented warming. Amongst the most important risks are the following:

- Sea level rise and loss of ice sheets. In the 20th Century global sea level increased by 15 - 20 cm. Currently sea level is rising at 3 cm/decade, faster than projected in the model scenarios of the IPCC. Future rise by 2100 will likely by less than one meter, but even if warming is stopped at 3 ºC, sea level will probably keep rising by several meters in subsequent centuries in a delayed response (Fig. 3). Coastal cities and low-lying islands are at risk. What is now a once-in-a-century extreme flood in New York City (with major damage, including flooded subway stations) would statistically occur about every 3 years if sea level were just 1 meter higher.
- Loss of ecosystems and species. Global temperatures would reach a high never seen for millions of years, and the rise would be much too fast for many species to adapt. A large fraction of species - some studies suggest up to one third of species - could be doomed for extinction already by the year 2050. Life in the oceans is not only threatened by climate change but by the equally serious problem of the ongoing global ocean acidification, which is a direct chemical result of our CO2 emissions.

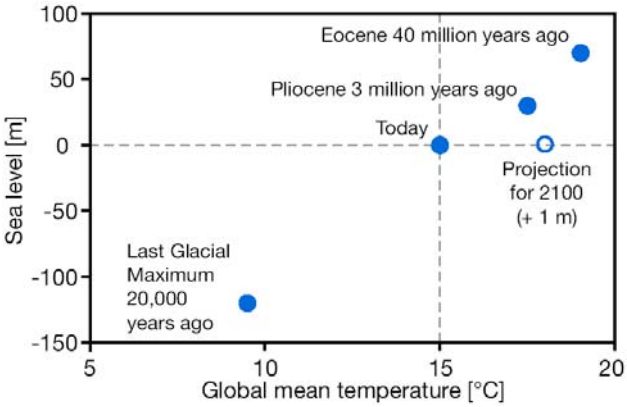

Fig. 3 Mean global temperature and sea level (relative to today's) at different times in Earth's history, with the projection for the year 2100 (1m above today's sea level). For the long term a much higher sea-level rise must be expected than that predicted for 2100. Source: after Archer, 2006.

- Risk of extreme events. In a warmer climate, the risk of extreme flooding events will increase as warmer air can hold more water (7% more for each ºC of warming). Droughts and forest fires are likely to increase in some regions, as is currently occurring in the Mediterranean region, Southern Africa and California. Hurricanes are expected to become more destructive. An increase in energy of hurricanes is suggested in response to rising sea surface temperatures by both models and data (Fig. 4). A number of recent studies has shown that the observed rise of sea surface temperatures in the relevant areas of the tropics is primarily due to global warming, not to a natural cycle.

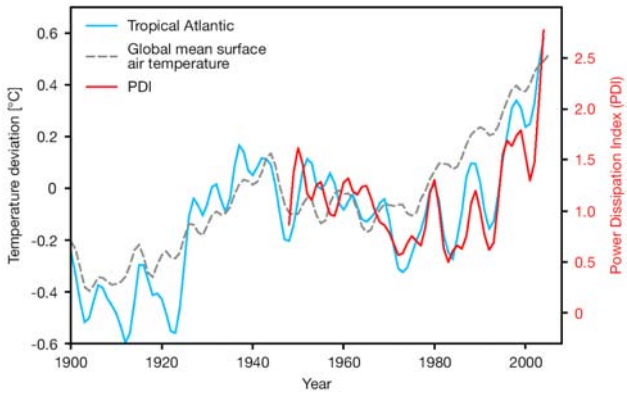

Fig. 4. Temporal development of the energy of tropical storms (Power Dissipation Index – PDI, red) and the average sea-surface temperature in the tropical Atlantic from August to October (blue). For comparison the evolution of the globally averaged near-surface air temperature is shown (dashed grey line). Source: after Emanuel, 2005

- Risk to water and food supplies. While the total global agricultural production may not decline in a warmer climate, many poorer and warmer countries can expect reductions in yields due to water shortages and weather extremes. The water supply of major cities like Lima is threatened when mountain glaciers disappear.

These are only examples – the exact consequences of such a major change in climate are difficult to predict, and surprises are likely. In some cases, impacts have already proven to be more rapid or severe than expected, like in case of the dramatic loss of summer sea ice in the Arctic Ocean. Ice extent in 2007 and 2008 was only about half of what it has been in the 1960s, ice thickness has decreased by 20-25% just since 2001, and in 2008 the North-East Passage and North-West Passage were both open for the first time in living memory.

How to avoid dangerous climate change

In the United Nations Framework Convention on Climate Change (UNFCCC) of 1992, almost all nations of the world have committed themselves to preventing a "dangerous interference" with the climate system. To avoid the most dangerous consequences of climate change, the European Union has decided to halt global warming below 2ºC above pre-industrial temperatures (EU limit, see Fig. 2), and many other countries have since endorsed this limit. To reach this goal, the greenhouse gas concentration in the atmosphere needs to be stabilised well below 450 ppm CO2-equivalent (possibly after some limited temporary overshooting of this value).

To achieve this, the global CO_2 emissions need to be at least halved by 2050, compared to the level of 1990. Carbon cycle feedbacks make this number uncertain, and the required reduction could turn out to be up to 70% (see Fig. 5).

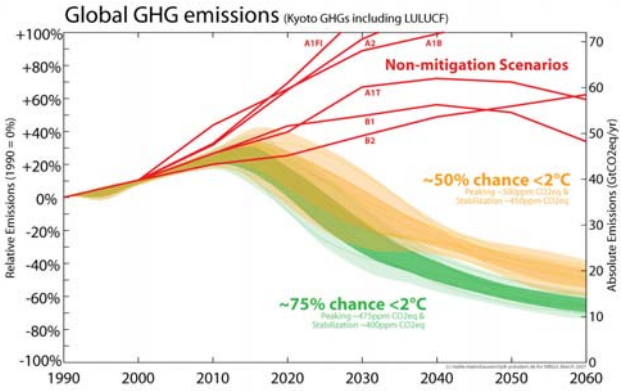

Fig. 5 Comparison of different emission scenarios for the period 1990 to 2060. Red curves are the well-known IPCC SRES scenarios without climate protection policies ("non-mitigation"). Yellow is a set of scenarios leading to a 50% chance to stay within the EU 2-degree policy limit. The green scenarios have a 75% chance of staying below this threshold. (Meinshausen 2007).

According to latest economic modeling results (see special issue of the *Energy Journal*, 2006, edited by O. Edenhofer et al., as well as the *Stern Review* published in November 2006), this can be achieved with minimal costs (less than 1% lower GDP by 2100) by induced technological innovation, including increased energy efficiency and renewable energy technologies (wind, biomass, solar). Detailed scenarios for the required energy transition have been worked out e.g. by the Advisory Council on Global Change of the German government (www.wbgu.de).

Sources

Archer, D., 2006: Global Warming: Understanding the Forecast. Blackwell, 256 Seiten

Emanuel, K., 2005: Increasing destructiveness of tropical cyclones over the past 30 years. Nature, 436, 686-688.

Edenhofer, O., Carraro, C., Koehler, J., Grubb, M. (eds) (2006): Endogenous Technological Change and the Economics of Atmospheric Stabilisation. A Special Issue of The Energy Journal, Vol. 27, International Association of Energy Economics, USA.

Mann, M. E., R. S. Bradley, and M. K. Hughes, 1999: Northern hemisphere temperatures during the past millennium. Geophysical Research Letters, 26, 759-762.

Moberg, A., D. M. Sonechkin, K. Holmgren, N. M. Datsenko, and W. Karlen, 2005: Highly variably Northern Hemisphere temperatures reconstructed from low- and high-resolution proxy data. Nature, 433, 613-617.

Oerlemans, J. H., 2005: Extracting a Climate Signal from 169 Glacier Records. Science, 308, 675-677.

WBGU 2003: „Beyond Kyoto" Special Report, (www.wbgu.de)

Figs. 3 and 4 are reproduced from the 2006 report *The Future Oceans - Warming Up, Rising High, Turning Sour* of the WBGU. (www.wbgu.de)

For more information, see the author's web page at www.ozean-klima.de and the climate change weblog www.realclimate.org

Life is a HoloDeck !

“Holography is a true intersection of art, science and technology”
(Prof. Stephen A. Benton, Holography Pioneer, 1941-2003)

Hi and welcome all to my lecture about 3D techniques, holography and spatial representation at the 25C3 ! It´s not only the 25th congress this year, sixty years ago in 1948 the inventor of holography Prof. Dennis Gabor published his papers in “Nature” which explained the holographic principle for the first time.

1. In the Beginning: Photography, Stereoscopy and Lenticulars

(The word "Photography" comes from the French *photographie* which is based on the Greek *φώς* (*phos*) "light" + *γραφίς* (*graphis*) "stylus", "paintbrush" or *γραφή* (*graphê*) "representation by means of lines" or "drawing", together meaning "drawing with light.")

In the fourth decade of the 19th century the first photographs were produced, soon after that the first stereoscopic photos occurred. The first 3D images of that period I have ever seen were made around 1845, showing posing girls in 3D. They could only be viewed with separators; today these glasses are called “anaglyphs”. Anaglyphs are still in use today, they can be made of different coloured filters, polarized filters (true colour, used for stereoscopic projections) or shutter glasses.

During the remaining 19th century we had the first stereoscopic cameras and even some first 3D film projectors occurred, when in the early 20th century F.E. Ives in 1908 filed a patent of “Parallax-Stereograms”. He had realized that he could watch stereographic images without any glasses when cutting them into small stripes and placing them under lenses (lenticulars), which directed the corresponding images to the left and right eye. This technique is called “Lenticular” and has made some good enhancements during the last decades. Today we can produce lenticulars in many sizes and types using these different techniques:

Lenticular origination techniques

Flip/change of 2 or more images
Animation of 6 or 8 up to a max. of 120 frames (short video sequences)
Zoom
Morph
Pseudo 3D (2D layers in different depth layers)
“Real” Photorealistic 3D (lenticulars with 3D views of an object)

When I say real 3D, it´s meant in regard to lenticulars, they give us a kind of spatial impression and depth, but with a quiet limited viewing angle compared to “real” holograms. Lenticulars are a lens-variation of stereograms but **not** a hologram.

Life is a HoloDeck !

“Holography is a true intersection of art, science and technology”
(Prof. Stephen A. Benton, Holography Pioneer, 1941-2003)

An interesting article about stereoscopic 3D and digital projection systems of Michael Starks can be found here:
http://events.ccc.de/congress/2008/Fahrplan/attachments/1184_Digital%203D%20projection%207-7-08_Michael_Starks.pdf

2. Holography - The Early Days

Theoretical Invention of Holography by “Coincidence”
Dennis Gabor–Emmet Leith/Juris Upatnieks–Yury Denisyuk-Stephen Benton

(**Holography** (from the Greek *ὅλος-hólos* = whole + *γραφή-grafē* = writing, drawing)

In the nineteen twenties the Hungarian physicist Dennis Gabor (1900–1979) studied at the Technische Hochschule Charlottenburg (1920-1927) and made his Dr. Ing. in electrical engineering in 1927. He often also stayed at the University of Berlin and met physicists like Einstein, Planck, Nernst and v. Laue.
From 1927 on he worked at Siemens and Halske as a research engineer here in Berlin. In 1933, when the Nazis took over the power, he fled and escaped to England. After the Second World War in 1947 he still worked at the British Thomson-Houston Company Labs in Rugby, UK. (1933-1948). At that time he researched on the improvement of the resolving power of electron microscopes.

Explanation:
The electron microscope at that time had a hundredfold better resolving power over the finest light microscopes, yet it still fell short of allowing scientists to "see" atomic lattices, since the resolution of the electron microscope had physical limitations.

The image was distorted in two ways: Fuzziness (as if one's camera were out of focus) and sphericity (as though one were looking through a raindrop). If one improved the former, the latter worsened, and vice-versa.

In 1947 a brilliant solution occurred to Gabor. What if one were to use the diffraction pattern (the fuzziness) in a way which provided one with all the information about the atomic lattice. That is, why not take an unclear electron picture, then clarify that picture by optical means. This was the genesis of holography. Gabor proposed to take a beam of light and split it in two, sending one beam to an object, the other to a mirror. Both would initially have the same wavelength and be in phase (coherent), but upon reflection from the object and the mirror back to the photographic plate, interference would be set up. Imagine ocean waves rolling in upon a long, sandy beach, one following another. Imagine them all equal in size, intensity, and timing. Now imagine you could split the beach in two, with two sets of waves coming in upon two different beaches. Tilt these two at an angle of your own choosing,

Life is a HoloDeck !

“Holography is a true intersection of art, science and technology”
(Prof. Stephen A. Benton, Holography Pioneer, 1941-2003)

superimpose them, and imagine the interference the waves would create for each other. This interference would not be completely chaotic, but would actually follow a pattern. From this "diffraction" pattern, one could reconstruct the initial waves. Now vary these initial waves in size, intensity, and timing (which might be imagined as due to different weather conditions out at sea). The diffraction pattern would differ correspondingly, and even the weather conditions might be hypothetically reconstructed. This is what Gabor wished to do with electron beams. The beam from the mirror would be unchanged, but the beam reflected from the object would contain all the irregularities imposed upon it by that object. Upon their meeting at the photographic plate, the two beams would be generally incoherent, and an interference pattern would occur. This interference could then be captured upon film, and if light were then shone through this film, it would take on the interference pattern and produce an image capable of three-dimensional reconstruction.

Gabor worked out the basic technique by using conventional, filtered light sources, not electron beams. The mercury lamp and pinhole were utilized to form the first, imprecise holograms. But because even this light was too diffuse, holography did not become commercially feasible until 1960, with the development of the laser, which amplifies the intensity of light waves. Nevertheless, Gabor demonstrated mathematically that holography would work even with electron beams--just as his experiments showed it worked with ordinary light. The major practical problem remaining with the electron microscope prior to 1960, however, was not left unchallenged by Gabor--this was the double image incidentally obtained by the holographic process. Gabor was able to use the very defect of electron lenses--spherical aberration--to remove the second image.

Gabor published the principle of holography and the results of his experiments in Nature (1948), Proceedings of the Royal Society (1949), and Proceedings of the Physical Society (1951). This work earned him in 1948 a position on the staff of the Imperial College of Science and Technology, London. In 1958 he was promoted to professor of applied electron physics, and he held that post until his retirement in 1967. His other work consisted of research on high-speed oscilloscopes, communication theory, physical optics, and television, and he was awarded more than 100 patents. Yet Gabor was not the pure scientist or isolated inventor; many of his popular works addressed the social implications of technological advance, and he remained suspicious of assumptions of inevitable technological progress, nothing the social problems it could not solve as well as the ones it caused.

Gabor received many honours. In 1956 he was nominated to the Royal Society; he was made an honorary member of the Hungarian Academy of Scientists; and in 1971 he received the Nobel Physics Prize for his holographic work. He died in London on February 8, 1979.
Modified from: http://www.answers.com/topic/dennis-gabor

Life is a HoloDeck !

“Holography is a true intersection of art, science and technology”
(Prof. Stephen A. Benton, Holography Pioneer, 1941-2003)

Extract:

Gabor was not capable to verify his theoretical forecast of the production of holograms at that time, because the laser (strong coherent light) was not invented yet. He used high pressure mercury lamps instead and so called In-line holography with just one single beam which interferes with itself. It´s different to the Leith/ Upatnieks holographic principle of using coherent laser light, splitting it into two half beams (object and reference beam) and then illuminate an object with the object beam thus creating the interference pattern in the exposed holographic plate. By doing so we **store the light reflecting characteristics of the object**, say the amplitude **and** phase of the light wave front being reflected from the object. (contrary to a photograph which only stores the intensity (amplitude) of the reflected light of an object).

Hologram Master:

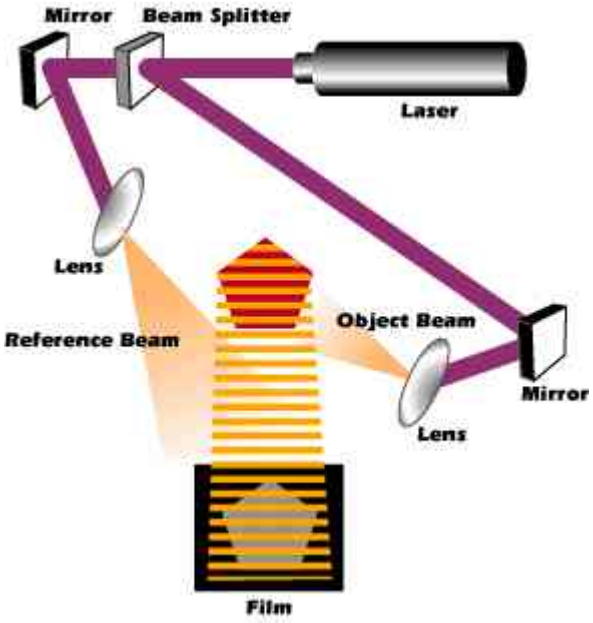

The laser beam is divided into 2 half beams with a beam splitter (50% mirror@45°), thus creating an object and a reference beam. The reference beam remains unchanged (as reference), while the object beam is reflected by the object and superimposes the reference beam at a special angle in the holographic plate/film. They generate an interference pattern containing the light reflecting characteristics (3D information) of the object in highest resolution. When a copy of the hologram is developed and viewable under white light, the white light reconstructs the optical light reflection characteristics of the object, just like looking through the master plate “window” directly at the object.

Gabor himself used the In-Line master set up, which is even simpler but nevertheless works perfectly too. Instead of using 2 half-beams the In-Line set

up consist of just one beam, which is reflected and interferes with itself (simplified).

Gabor invented several other techniques just by „coincidence“, not only his social ambitions and scepticism against uncontrolled abuse of technology and progress make it worth to study some of his work. You can have a look at his Lecture, Autobiography and Nobel Prize speech of 1971 here:

http://nobelprize.org/nobel_prizes/physics/laureates/1971/gabor-lecture.pdf

http://nobelprize.org/nobel_prizes/physics/laureates/1971/gabor-autobio.html

http://nobelprize.org/nobel_prizes/physics/laureates/1971/gabor-speech.html

The different principles, types and processes of hologram production are often found in the web, for instance at www.holography.ru , a page of Russian holographers who offer courses, their excellent and detailed hands-on guide can be found here: http://www.holography.ru/techeng.htm

3. **Laser** - **L**ight **A**mplification by **S**timulated **E**mission of **R**adiation

In 1960 Theodore Maiman (1927- 2007) invented the first ruby-laser. In 1962 two scientists at the MIT, Emmeth Leith and his assistant Juris Upatnieks remembered Gabor´s forecast of holography in 1948 by using a stable coherent light, which now was available. They decided to try out the principle “by pure curiosity” and they used the 2 half-beam set up with object and reference beam for the first time.

It obviously worked and they created the first hologram in the world, showing a model of a Train and Bird:

Life is a HoloDeck !

“Holography is a true intersection of art, science and technology”
(Prof. Stephen A. Benton, Holography Pioneer, 1941-2003)

Now the theoretical invention of holography in 1948 by Dennis Gabor was proofed in practice for the first time. The holograms at that time were only viewable under laser light of the same wavelength used during the exposure; laser masters today still are viewable under laser light only.

Since 1962 the Russian physicist Yury Denisyuk (1927-2006) developed another improvement by combining Lippmann´s work in true-colour photography with holography. The first two-colour reflecting hologram which could be illuminated with white light was produced in 1965 by G. W. Stroke and A. Labeyrie. Denisyuk invented the Lippmann-Denisyuk-Stroke reflection hologram using a different geometry: The object and reference beams meet in the holographic plate from opposite directions and create wave patterns in the depth of the emulsion. Now reflection holograms viewable under white light became commonly available.

Embossed Holograms

Another important improvement in holography was initiated in 1968 by Stephen Benton (1941-2003). Benton developed the first embossed rainbow or Benton-Hologram. This type of hologram reduces the vertical parallax by using a slit in the set up of copying the hologram. It creates the typical spectral rainbow effect well-known from embossed holograms (on plastic-cards, labels or money). This technology provided the opportunity to mass produce holograms in huge amounts. It´s exposed using a special emulsion, the photo resist. When the hologram is exposed and developed in the photo resist plate, it´s surface can be silvered which makes it electrically conductive. This allows us to bath the plate in a galvanic electroforming tank to create a nickel copy with a thickness between 100 and 300 µm. It can be used as a printing plate (nickel-shim) for embossed holograms. This is our grandmother-shim (positive), we can pull a mother shim through electroforming again (negative) and finally grow the production daughters positive again to emboss into the film (since the production daughters have a limited life time, you can always pull another one from the mother; holographic galvanic is quiet similar to CD/DVD electroforming but at a higher resolution).
The daughter nickel shims are finally mounted on a rotary embossing cylinder (roller) in a holographic embossing machine. Under high pressure and heat the microstructure is transferred into a thin lacquer of ~ 3 µm on a film made of PET, PE, OPP, BOPP at al. Behind the lacquer is another even thinner layer of some nanometre aluminium atoms which reflects the white light. The thickness of the film itself is usually between 12 µm (thin hot-stamping foil) over 25, 30 to 50 µm (thick label/sticker foil).

The actual physical / chemical process of creating a hologram in a master lab is quiet similar to photography except to the absolute sensitivity of the recording system of a hologram. We expose the master hologram plate (H1), then put it into a chemical developer, fix and finally bleach the plate. Since we are working with interference structures in the nano dimension (normally within the viewable light spectrum between UV and IR / ~ 400–800 nm), we

Life is a HoloDeck !

“Holography is a true intersection of art, science and technology”
(Prof. Stephen A. Benton, Holography Pioneer, 1941-2003)

have to avoid any vibrations within the recording geometry larger than ¼ of the wavelength of the laser used.
Imagine we use an argon or crypton laser in green, it has it´s most effective wavelength at 488 nm. Having a vibration of 122 nm inside of the system will completely destroy all interference patterns and make the hologram invisible ! A vibration of 122 nm in a lab recording system is similar to maybe one degree C temperature difference in the room or if you just say a word, so during exposure you better close the door and shut up.

For that reason we build and use heavy special master tables to do such exposures, normally we set up the optics and mirrors first, then leave the lab to have a coffee or something (settling time needed to damp the whole system down from any vibration) and finally making the exposure preferable remote controlled by a switch/shutter from another room.
The exposure time depends on the used emulsion, the power and stability of the laser and size of the Hologram. It usually lasts between some seconds and a few minutes.
(Except in Pulse Laser Portrait Holography: to solve the problem of a person moving during the exposure huge accumulators are charged up to some Megawatts first and then flash a laser beam, widely opened by optics of course, for app. 50 nanoseconds)

The exposed and developed H1 is only viewable under laser light of the same wavelength used during the exposure, but it certainly is the most impressing and effective Hologram we can produce and observe. It has a viewable depth of 10 or even more meters with full horizontal and vertical parallax !

The recording emulsions have their highest sensitivity at the laser wavelength used, in the beginning monochromatic only, meanwhile polychromatic to expose RGB colour holograms using red, blue and green lasers. The corn of such emulsions is less than 8 nm today, giving us resolutions of 20-80 nm or > 15.000 lines per mm in the H1, while the resolution of a typical rainbow embossed hologram decreases to 1500 lines/mm “only”.

Because of that high resolution holography is ideal for precise measuring made in interferometry and the first holographic data storage disks are also now available:
http://www.inphase-technologies.com/

Here is a brief glossary of basics and related issues, if you have additional questions we can discuss those and I will try to answer within the remaining time, I´d like to thank you for your interest and thanks to Wetterfrosch too, who has helped getting this lecture done and for layouting the 25c3 space lenticular cards, we hope you enjoy and spread them all over the world, thanks and good nite !

Homemade Mask hologram:
http://www.youtube.com/watch?v=mqV56adpBpQ

Life is a HoloDeck !

“Holography is a true intersection of art, science and technology”
(Prof. Stephen A. Benton, Holography Pioneer, 1941-2003)

Glossary:

H1 – Laser Transmission Master Hologram
Classical Transmission Holograms made on a mastering-table

Denisyuk / Reflection Hologram
White light copies of H1´s - Denisyuk / Reflection Holograms

Rainbow / Benton Hologram
Photoresist mastering and Electroforming to produce embossed Holograms

Dot-Matrix Origination
Computer generated Pixel Masters for embossed Holograms

Micro- + Nanotext / Micropoint
Very small text or code within a Hologram

"Hidden" information
Invisible text or code becoming visible when illuminated with a laser

Multiplex / Integral-Hologram
An Integral Hologram is a combination of stereoscopy and holography which is made out of many single images. They can be recorded with a moving camera, thus getting a spatial information and movement of the object or person. The frames are copied vertically into the hologram with the laser. Any kind of computer generated object or image can also be taken.

Pulse Laser or Portrait Holograms
Very short pulsed laser flash to produce holograms of humans, animals etc.

True-Colour and Pixel Reflection Hologram Mastering
True-Colour Holograms are exposed in emulsions sensible in RGB with RGB laser systems.
Pixel Reflection Holograms are exposed using a special printer which generates millions of little pixel holograms of ~ 0.5 - 1.5 mm size. Can be 2 or true colour and also be generated from rendered virtual data or video.

Holographic Special Machinery:

Recombining Systems
Electroforming Baths
Different types and sizes of conventional holographic embossing machines - Soft- and Hard Embosser
New generation manufacturing technology and equipment - UV/Electron Beam Casting
Converting Equipment

Privacy in the semantic web– Social Networks based on XMPP

Jan Torben Heuer <jan.heuer@uni-muenster.de>

December 13, 2008

1 Introduction

In the last years the static web has moved towards a socially interactive web. Since 2005 we often refer to this as the *web2.0* [1]. People collaboratively write articles in online encyclopedias like Wikipedia[1] or self-portray themselves with profiles in social networks like Myspace[2]. Within a social network, members can link with each other in order to create a personal network of friends. Often, the number of friends is a kind of "social status" and displayed on a person's profile page. According to [2] this community aspect is one of the reasons why social networks attract so many users.

In current social network applications we see a **lack of privacy**. If private data is shared with others across a web based social network then our privacy can be affected on different levels. The minimum privacy requirement is that private data must not be visible to everyone. Therefore most services provide a simple access control for content. Flickr for instance supports private or family photo albums. Recently some security issues arose in social network which allowed users to gain private data of other users. In [3] has been shown how to get limited control over the computer of Facebook users in order to perform a denial-of-service attack against another host. The authors also explained how malicious applications can access a private user-profile in order to copy the private data to a remote server. So another privacy requirement is that data is kept secure – a trust that our application must also gain, of course. Most important is the trust in the service provider regarding what he actually *does* with private information. Besides explicitly provided information like preferences there are also implicit information from user tracking – visited profiles, groups or advertisements. Mostly such data are used for market analysis and user specific advertising. Our privacy is the price we have to pay in order to use such free services. But the price may be too high. The issue is not the advertising but the continuous collecting of data. Information that have been gathered once cannot be reverted. The membership in a discussion group about certain diseases for example can prevent someone from getting hired if the employer gets this knowledge.

[1] `http://www.wikipedia.org`
[2] `http://www.myspace.com`

Moreover, such information can also cause disadvantages for later generations in case of genetic diseases. We are the first generation in the "information century" and long term impacts of the ongoing user profiling in social networks are not foreseen, yet. Social networks are a phenomenon which attracts a lot of users, especially users who do not have the technical knowledge to understand what happens in the background. Or what is possible with current user tracking software.

Additionally most content in social networks is kept in proprietary formats. There is no content integration between the existing independent social networks. All have their own community and neither user-accounts nor friendships can be shared between the social networks. Ideally data should be exchanged in an open and extensible format such as the Web Ontology Language. And here we we see the **conflict**. We cannot ensure privacy and ask for interoperability at the same time. On the one hand, an open exchange of content between different sites makes it easier to collect information about a user. An exchange of user profiles between social networks simplifies the aggregation and reasoning of hidden information. On the other hand, to ensure privacy protection the different social networks would have to restrict public access and use proprietary instead of open formats. This way, the vision of an open, interconnected and interoperable semantic web[4] won't come true.

Regarding this conflict between interoperability and privacy we propose an alternative to traditional server-based social networks. In this paper we describe a distributed social network architecture and compare it to currently used server based applications. Like in traditional social networks users create an account which identifies them in the network. The users can create personal profiles and link to others in the network. Our architecture does not need a central server which stores private data of the users because personal data are shared directly with linked friends. Our application therefore does not have the privacy issues of social networks that we identified above.

2 A network architecture based on XMPP

In current web-based applications users connect to and store their (private) information on central servers. Figure 1(a) illustrates the principle design. The opposite of a client-to-server network is a *peer-to-peer* (P2P) network. In P2P networks, each peers are similar, which means that they can act as a server (data provider) and client (data consumer) at the same time. There is no central instance that routes the traffic or knows about all connected peers. Peer-to-peer technologies [5] are mainly used for anonymous filesharing where one peer tries to connect to as many other peers as possible. In contrast, the authors define a network where peers only communicate with manually selected peers as a *friend-to-friend* (*F2F*) network. We use a F2F network for our application because we only want communication between clients that *trust* each other. The authors additionally point out that a serverless network is not able to guarantee connectivity because clients are often located in a local network that is connected to the Internet through a router. In such configurations the router acts as a packet filter (often called *firewall*) with NAT (Network Address Translation). Two peers that are both located behind a NAT router

cannot directly contact each other because connections must always be initiated from inside the local network to the outside (i.e. the Internet). For that reason a public available message relay is necessary where clients can initially connect to. The *extensible messaging and presence protocol XMPP* ([6]) is such a messaging network architecture for client-to-client communications (see figure 1(b)).

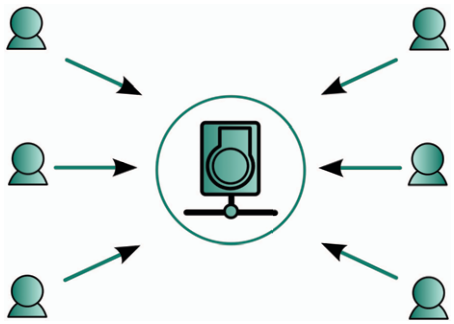

(a) Client-server architecture: a server stores the client's data in its database. Incomming queries are evaluated by the server.

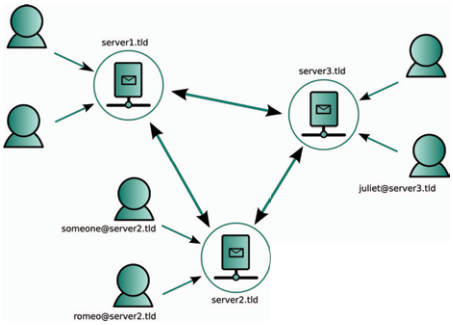

(b) Messaging network: clients store their data, the servers are only responsible for communication. Clients directly query each other for information.

Figure 1: Differences in communication and data location

A XMPP network consists in a set of independent servers. The servers needn't know each other initially and anyone can set up his own XMPP server. Users register at their preferred server and get an user–id which is similar to an e–mail address: `user@server.tld`. Two clients which exchanged their addresses (and authorized each other for communication) can then send messages to each other. The servers only handle the message routing and presence information (if a user is currently online) but do not need to know about the concrete information that are exchanged. At the beginning we argued that server based social networks in the web cannot ensure privacy. Although we also rely on servers in our architecture their role is different. Servers are only responsible for transferring messages but needn't know about the message contents.

In order to provide a private message exchange for clients we use the public-private-key encryption PGP. Messages can be encrypted to protect them against eavesdroppers and messages can be signed to prevent forgery. This technology is already used in the XMPP context for encrypting instant messaging but it isn't widely used, yet. There are two steps which are important in order to establish a really secure connection: At first, a pair of a public and private key has to be created. This can be done automatically by our application when a new XMPP account is created. Second, public keys have to be send to friends and their public keys must be verified. If someone pretends to be a certain person and introduces himself with fake name there won't be a possibility to automatically verify the identity. Like in other social networks anyone can sign up with any name. Unless we have a central authority which maintains identities - and we clearly don't want to have such an institution - people have to manually identify each other. That means, they have to verify that a public key is really owned by the the person it

claims to. If this check isn't done carefully the future communication cannot be assumed to be secure. Typically, the verification is perfomed by compareing the fingerprint of the public key either by telephone or face-to-face. Our application can only support the users by i.e. asking for a friends fingerprint but we do not want to include any automation.

3 Data storage and exchange

In the article "The Semantic Web" [4] the authors described their vision about the future of the web. While the current web is for the people, the semantic web shall allow agents to read, understand and process the available information in order to support people in information retrieval and decision making. Although we're still far away from this vision becoming true, a lot of work has been done in recent years to archive the goal. The World-wide-web Consortium W3C developed the resource description framework RDF which is a general tool to model information in a structured way. Information are written as statements or triples of `{Subject, Predicate, Object}`, for example "user1 hasName Bob". The triples can be read by agents and processed by logic-interpreters (reasoners). Other popular data models are relational databases or XML documents. While relational databases store data in a tabular structures or XML documents have a tree structure, triples are graphs. RDF graphs for example do not have a root and can contain circular referenes. Each statement's subject can also be an object in another statement, for example "user1 knows user2".

Ideally, information are stored in a semantic format, distributed over the web and linked to each other. Today, the web concentrates at some websites. Most web based social networks do not provide an export of user data or have their own proprietary format or API. Although this formats are often structured with open standards like XML their content isn't interoperable. Each social networks defines its own structure. Our goal however is to directly share information between clients without the need to convert them manually. Also we don't want to bound the user to our application just because he doesn't get his data out of it. Therefore we use existing RDF schemas to model the user's information. The main schema is the FOAF vocabulary [7] which is used to share personal data like name, e-mail address and links to other friends. A graph of the vocabulary is shown in figure 2.

As an initial example we implemented a social semantic tagging application. Users can create their own personal profile and bookmark websites. The bookmarks can be annotated with keywords, so called *tags*. For this *tagging* already exists a RDF schema, the TagOntology [8]. Further extension can define their own schemas and connect it to the existing ones. If one wants to query information from his friends, i.e. their names and e-mail addresses or bookmarks for a certain tag he can use SPARQL [9] queries. The syntax is similar to SQL but for RDF statements instead of relational databases. This generic query approach allows for an easy retrieval of any resource. Of course, each user can define which personal data in his database are available to whom.

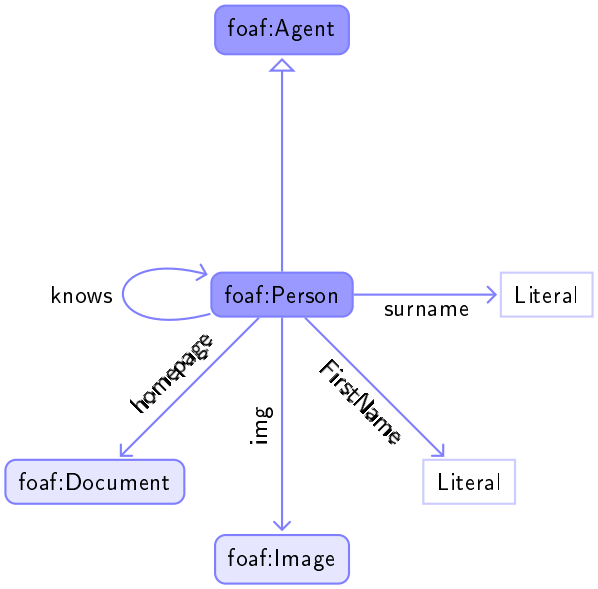

Figure 2: The classes Agent and Person from the FOAF vocabulary and some properties as graph. Person instances can reference to other Persons with the `knows` property. There are much more properties defined, i.e. to describe Jabber or OpenID accounts.

4 Application

We developed a first desktop application in Java for sharing bookmarks with friends. The latest Java release (Java 6) has a much better cross-platform desktop support than the releases before, for example "System Tray Icons" and access to the system–default web browser. Another reason for Java 6 was the Java Webstart technology which allows for automatic update detection on startup. The application can directly be started from the website `http://www.pace-project.org` which has been created for this project. However we want to stress out that our application is only one possible implementation of the concept. It might also be interesting to include the functionality in exiting instant messaging clients. The application can also run as a small server with a web frontend for the local browser so that it looks familiar to other web-based social networks. The main reason for an implementation in Java was that most functionalities are already provided by other projects. The XMPP protocol was implemented by the "Smack API" from igniterealtime[3]. PGP security functionalities are provided by the "Bouncy Castle Crypto APIs"[4]. For the database engine we use "Sesame" from Aduana[5]. It is a full-featured RDF store that directly supports SPARQL queries. In order to let people use their existing bookmarks we wrote two import plug-ins for Delicious and Bibsonomy. If a

[3] `http://www.igniterealtime.org/projects/smack/index.jsp`
[4] `http://www.bouncycastle.org/java.html`
[5] `http://www.openrdf.org/`

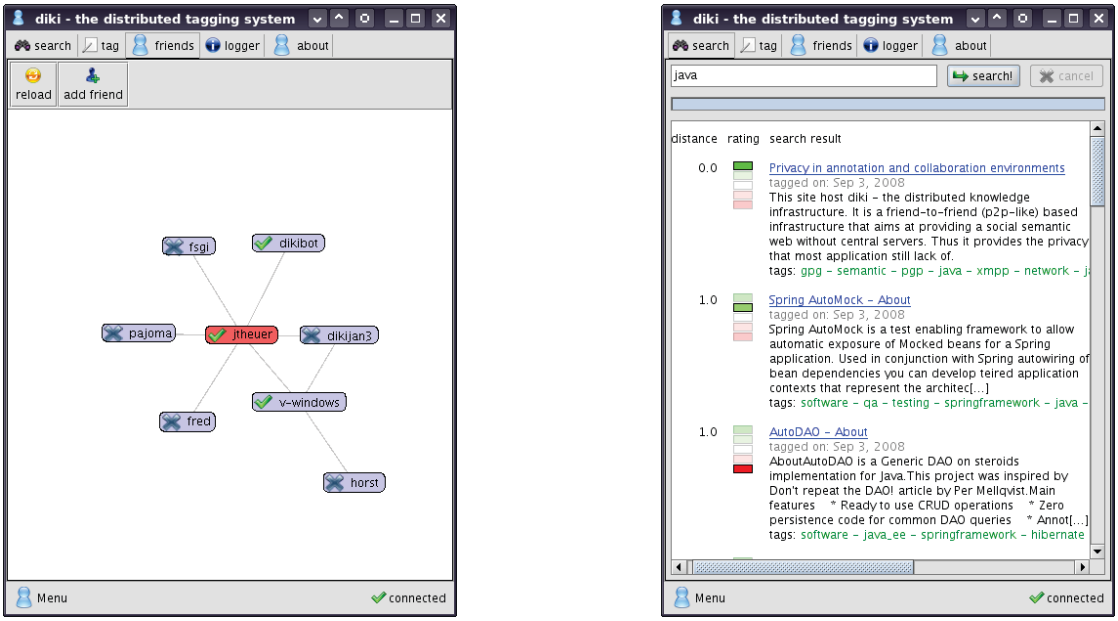

Figure 3: Two screenshots of the application. The social friends-graph (left) and a search result for the tag *java* (right).

user already has an account for the services he can link his bookmarks in to application.

Other locally running programs can interact with our application in order to trigger certain actions. The GUI provides access to its API with the integrated HTTP server jetty[6]. By clicking the "new Bookmark" button in Firefox for example, a HTTP request to our application is triggered: `http://localhost:30013/api/newBookmark?url=...`. The application opens the "new bookmark" window of our application and the currently displayed URL is filled in and can be annotated with tags. A similar plug-in also exists for the KDE web- and file browser "Konqueror". The HTTP API can also provide additional services. Many people for example like to aggregate news in an RSS reader (RSS is a good example for a widely used RDF-based format). Sites like Delicious provide a "hotlist" which are the currently most often bookmarked sites. Our application can generate RSS feeds of the pages that have recently been bookmarked by friends.

5 Challenges and future work

Our application will only work if sufficient clients are online, like other client-based applications (i.e. for filesharing or instant messaging). Social networks however will only be attractive if a certain amount of interesting information are available (profiles of friends or bookmarks). Therefore we need a mechanism to cache data in the network. One possibility is that users can cache profiles of friends and distribute them while they

[6] `http://www.mortbay.org/jetty/`

are offline. But we have to make sure that neither the privacy of friends is affected (by giving information to the wrong people) nor that conflicts because of different versions occur. Another future goal of our application came up when we integrated PGP security. Currently, PGP public-keys are shared over keyservers. All data from this servers can be downloaded by everyone. There are for example keys, e-mail addresses and a set of users who signed the key. The singers are possible friends of the keyowner and therefore we can say that the keyserver's content is one of the first social networks[10]. But some people don't want to expose their information and therefore avoid the keyservers. Nevertheless keyservers are a convenient way to find and exchange public keys. We propose to use our architecture as a keyserver replacement. PGP keys for E-Mail communication are added to a user's FOAF profile and can be queried by friends only. The application provides a standard PGP keyserver interface on its default port for the localhost. Other applications like e-mail programs can then access keys of friends like before because internally the key search is translated to a SPARQL query in our application.

References

[1] Tim O'Reilly. What is web 2.0: Design patterns and business models for the next generation of software. *Social Science Research Network Working Paper Series*, September 2005.

[2] Adam Mathes. Folksonomies – Cooperative Classification and Communication Through Shared Metadata. `http://www.adammathes.com/academic/computer-mediated-communication/folksonomies.html`, December 2004.

[3] Elias Athanasopoulos, A. Makridakis, Spyros Antonatos, D. Antoniades, Sotiris Ioannidis, Kostas G. Anagnostakis, and Evangelos P. Markatos. Antisocial networks: Turning a social network into a botnet. In Tzong-Chen Wu, Chin-Laung Lei, Vincent Rijmen, and Der-Tsai Lee, editors, *ISC*, volume 5222 of *Lecture Notes in Computer Science*, pages 146–160. Springer, 2008.

[4] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific Amercian*, May 2001.

[5] Michael Rogers and Saleem Bhatti. How to disappear completely: A survey of private peer-to-peer networks. In *1st International Workshop on Sustaining Privacy in Autonomous Collaborative Environments (SPACE 2007), Moncton, Canada*, July 2007.

[6] Peter Saint-Andre. Extensible messaging and presence protocol (xmpp): Core. Technical report, IETF proposed standard, RFC 3920, Oct 2004.

[7] Dan Brickley and Libby Miller. FOAF Vocabulary Specification. Namespace document, FOAF Project, September 2004.

[8] Richard Newman, Danny Ayers, and Seth Russell. Tag ontology. Technical report, December 2005.

[9] Eric Prud'hommeaux and Andy Seaborne. SPARQL query language for RDF. Technical report, World Wide Web Consortium, 2005.

[10] Robert H. Warren, Dana Wilkinson, and Mike Warnecke. *Empirical Analysis of a Dynamic Social Network Built from PGP Keyrings*, volume 4503/2007 of *Lecture Notes in Computer Science*, pages 158–171. Springer Berlin / Heidelberg, April 2008.

Society

Social Movements Against the Global Security Architecture!

A Critique of the Militarisation of Social Conflict and the Securitisation of Everyday Life

By Gipfelsoli (full text on http://www.gipfelsoli.org/Gipfelsoli/5575.html)

Recent unrest due to food price hikes, protests against rising energy costs, visions and realities of a climate crisis and growing concerns over scarce resources, in conjunction with the continued turmoil of financial markets, are creating a sense of insecurity for a neoliberal regime in severe crisis. The G8 states and their allies are seeking to contain these conflicts and the evident accumulation crisis of the global economy through market-orientated solutions in order to restore economic growth whilst calls for more state intervention in the regulation of financial markets are rife. At the same time, the 'war on terror' serves to justify ever-more militarisation of all spheres of life. Wars are waged to secure new markets, transport routes and resources. New techniques of governance are emerging within a logic of waging war against who- or whatever cannot be made profitable.
In 2009, a number of security policy changes whose consequences are as yet unclear, are planned for the EU.

Under the banner of 'civil-military cooperation', internal and external security are to be merged into a 'comprehensive' and supra-national 'security architecture'. With a view to the US ministry, the 'Department of Homeland Security', founded after September 11th 2001 and comprising governmental, business and research organisations, EU security authorities are pushing for similar policy approaches for the European Union. 'Homeland Security' is to form the basis of the global security architecture of the most dominant states and richest economies, incorporating supranational institutions and multilateral agreements.

In a paper titled, 'Freedom, Security, Privacy – European Home Affairs in an Open World', the 'Future Group', initiated by German Home Secretary Wolfgang Schäuble in 2007, demands are made for a profound change of course in EU Home Affairs towards 'Homeland Security' (although the precise term is not used). Europe is to take the lead in the response to "security, migration and technological challenges". Priorities are:

> *"Police cooperation, the fight against terrorism, management of missions in third states, migration and asylum as far as border management, civil protection, new technologies and information networks".*

Every five years the EU decides upon new guidelines for the 'internal security' of its member states. Following the Tampere Programme (1999-2004) and the Hague Programme (2004-2009), the current paradigm shift is to be implemented in the second half of 2009 during the Swedish presidency. Characteristic for this paradigm shift is an "early strategic diagnosis" to prepare for "threats" that are not yet reality but are imaginable. Through the use of "risk analysis" dangers are projected that 'necessitate' ever-more militarised internal security carried out in close cooperation with think tanks and civil defence authorities.

"Transforming the Data Tsunami into Intelligence"

The EU Strategy paper, 'Freedom, Security, Privacy – European Home Affairs in an Open World' [2]

The changes to Home Affairs policies stipulated in the Hague Programme have already been implemented by many member states: Harmonisation of terrorism legislation, data retention, development of existing databases with common access, cross-border police cooperation, for example during sporting events or mass political protests. 'Border management', finger-printing when applying for an EU visa, as of 2009 biometrical identification in new passports and ID cards, development of security research, cooperation in crime management, police missions abroad etc. On the level of the EU new institutions have been founded and existing ones have been granted more responsibilities.

Many of the regulations described above were announced as temporary measures in the 'war on terror' following 9/11. Today this state of exception has become the norm and is being further exacerbated. According to EU papers, other principle threats to the "European Model" are migration and "organised crime". In their paper, 'Freedom, Security, Privacy – European Home Affairs in an Open World', the 'Future Group' suggests three "horizontal demands" for European security and

proposes “the development of Europe’s position in a globalised world”:

- *“preserving the ‘European model’ in the area of European Home Affairs by balancing mobility, security and privacy*
- *coping with growing interdependence between internal and external security*
- *ensuring the best possible flow of data within European-wide information networks”*

The changes in European Home Affairs policy are directly connected to discussions regarding the EU Lissabon Treaty that seeks to create EU-wide organisations to strengthen the economic competitiveness of the EU. The “convergence principle” enables the dismantling of intra-state and legal “obstacles”, legislation is to be “harmonised” and “simplified”, equipment and personnel is to be “pooled”, education and training is to be standardised with the aim of “interoperationability” between existing systems. As with many areas within the EU, the decision-making structures for Justice and Home Affairs are not transparent.

The stated goal of the EU Treaty is to create a space of ‘Freedom, Security and the Rule of Law’. Since 2008, Jacques Barrot is the ‘Commissioner for Freedom, Security and the Rule of Law’. He took the place of Franco Frattini who became the ‘Forza Italia’ foreign minister of Italy within Berlusconi’s new cabinet. Together with the German Home Secretary Wolfgang Schäuble, Frattini played a central role in developing the current European security policy. In the area of inner security the relatively young EU institutions are to receive more operative competencies in order to keep social discontent within the EU under control: the European police academy CEPOL, the European Gendermerie EGF, the police agency Europol, the EU “Situation Centre” (SitCen) for the analysis of current situations and for the evaluation of intelligence and the border protection agency Frontex. They are to have access to all “relevant information”. Following Germany’s example, all member states should set up “anti-terrorism centres”, merge police and intelligence services and exchange information on an EU-wide scale (so long as the respective ‘national interests’ are not endangered by the exchange of intelligence). The intention is to create an overarching ’Committee for Internal Security” for all EU police authorities. This is a step towards a common “EU Home Affairs Ministry”, beyond the existing “Standing Committee on Internal Security” (COSI).

The coordination amongst EU security authorities is conducted by so-called “Liaison Officers” who already have a large number of responsibilities and competencies with a department in each member state. The ‘Future Group’ advises expanding and strengthening their network.
Following the logic of ‘Homeland Security’, internal security is not merely a matter of Home Affairs, but is a common effort comprising of policy-making, military forces, police, civil defence, the security industry and think tanks. The intention is to “blur” the boundaries between these authorities and recognise how they are “intrinsically linked“ due to the need for a “comprehensive global approach”.
Consequently, under the primacy of security, different policy areas are interwoven:

> *“The Group strongly advocates developing a holistic concept covering e.g. development, migration, security, economic, financial, trade and foreign policy aspects in this regard, allowing the European Union to play a responsible and credible role in international relations“*.

As a contribution to the ‘global security architecture’, internal security should also be organised between different states. The main focus is the USA with whom the EU already has a number of bilateral agreements with respect to data exchange, Europol, extradition, mutual assistance, passenger information, SWIFT transactions and container security.

> *“In the area of freedom, security and justice, actions and measures have to follow strict geographical prioritisation and political differentiation: the European Union first has to define its key strategic interests. […] At a second stage the European Union has to identify which third countries are of vital interest for cooperation.“*

“Geographical challenges” are “candidate countries”, the West Balkans, neighbouring countries of the EU, the Mediterranean region, Russia, Africa, Latin America, Afghanistan, Iraq and its neighbouring countries, China and India. The ‘Future Group’ concludes that “the blurriness of internal and external threats” and an “internationalisation of conflict resolution” makes “interventions outside of the EU“ necessary. The aggressive foreign policy of the EU is nothing new, but the fact that Home Affairs ministers now consider it one of the highest priorities, marks a new era.

The EU police mission EULEX deals with tasks such as fighting insurgency, protecting private property and maintaining public order. The common “European Gendarmerie Force” (EGF) which has its base in Vicenza in Italy, is to be incorporated more in foreign missions. Police deployments are understood as “civil interventions”. In future, all forces operating abroad (military, police, diplomatic

services, development ministries, civil defence, “institutions for the rule of law”) are to have recourse to knowledge and information they need in the planning stages already, and operate together in ‘Mission Situation Centres’, making their information available to other EU authorities. Potential operating areas are:

> *“Institution-building, rule-of-law missions, election monitoring, democratisation, civil society and humanitarian aid. […] The vast variety of threats ranges from war situations to terrorist attacks, organised crime, violent demonstrations, natural or man-made disasters and usual police tasks.“*

Within member states, but also on an EU level, new decentralised “competency centres” will be developed in order to bundle common tasks. The increasing interweaving of police and intelligence work necessitates common surveillance centres for the whole range of telecommunication surveillance. At the external borders of the EU, common “Police and Customs Cooperation Centres” (PCCC) are to be installed.
EU ministers have declared migration an “inherent phenomenon in our increasingly globalised societies and economies“. The economic aspects of migration are at the forefront of their considerations. Demographic developments are registered with concern, whilst prognoses are made regarding the increase in much-needed labour migration. Thus, so-called “legal migration” is to be further strengthened in order to provide labour for the EU labour market. The relationship between “supply and demand” between the EU and workers from “third countries” will form a back-door option to ‘return’ migrants in accordance with demands of the labour market:

> *“The Group suggests that Member States should fully exploit all possibilities of intra-European economic migration.“*

In third countries more EU migration authorities will be installed,

“with responsibility for advising on visa and related questions and recruiting possible immigrants.“

The EU is to implement an “Entry-exit” system similar to the US ESTA system (“Electronic System for Travel Authorisation”) that will be operationalised on January 12th 2009. The US system is designed for tourists and business travellers. Mobility is a central issue for labour market policy and tourism and thus has a security dimension:

> *“If citizens do not feel secure, then it is highly likely that they would not wish to travel at all”.*

Having an EU travel document should facilitate the crossing of borders:

> *“A one-stop approach integrating all checks and controls carried out for different purposes, i.e. relating to persons, goods, veterinary and phyto-sanitary, pollution, terrorism and organised crime.“*

Here, new ‘border management’ technologies are to be introduced (e.g. biometrics, x-ray technologies, RFID chips). Further to the harmonisation of asylum laws, the “fight against illegal migration” requires an “effective European return policy”. The European border control system (EUROSUR) is to be expanded,

“to reduce the number of illegal immigrants entering the European Union by improving situational awareness at external borders and increasing the reaction capability of information and border control authorities.“ [3]

For this, existing institutions and programmes will be further networked. The focal point for this will be the “border protection agency Frontex” in the “fight against migration”, “organised crime”, drug trafficking and terrorism.

> *“The success of Frontex missions to date is undermined by the lack of precise legal provisions on, for example, the regime governing Frontex measures with regard to e.g. sovereign action executed by national ships or planes and responsibilities for refugees, asylum seekers and castaways. Therefore, priority should be given to the development of such common rules.“*

Accordingly, member states should do more to grant Frontex more responsibility in short-term “missions”, in setting up regional departments and providing technical assistance and other materials. Frontex should not only train national border protection troops, but also inspect and evaluate them. More deportations (“return flights”) are to be carried out under the independent “initiative, organisation and coordination” of Frontex. The aim is to develop a common “corporate identity” of all EU border troops as “European border guards”. EU border troops should also operate outside of the EU, for example between Lybia, Niger and Chad. At sea their responsibility should be expanded to the “territorial waters of affected third countries”.

Of great concern to Home Affairs ministers is the standardisation of security technologies. Civil and

military research should be further merged in the “European Security Research Programme” (ESRP, the ESRP alone has a budget of 1.4 billion Euros for 2007-2013). The links between national and supra-national authorities, private businesses and think tanks are evident in the example of how Germany contributes to the ESRP; through representatives of the BKA (Federal criminal investigation authority), the ‘Fraunhofer Gesellschaft’ (a leading scientific think tank), the arms companies Siemens, Diehl and EADS.

European police authorities are irritated by data protection, the increasing use of encryption software and encrypted telecommunication (PGP, Skype). The proposal is to develop standards for the future that make it easier to carry out bugging operations. Also in the area of video surveillance the intention is to harmonise the systems in order to diminish technical problems that stand in the way of common access to biometric data for example. More research is to be undertaken regarding police deployment of “un-manned systems” (so-called “Unmanned air vehicles” (UAV), remote controlled “drones” with cameras installed). A number of police forces in Europe are testing the use of UAVs within police operations more generally. The use of UAVs in Switzerland has already led to arrests of migrants along the “green border”.

Data bases and new technical developments play a central role in the redesign of EU Home Affairs. This logic ensures that security and individual rights can only flourish in an “atmosphere of collective security”, as the former EU Commissioner Frattini explains. His words are echoed in the strategy paper; New technologies and common databases,

> *“ensure more security for citizens and at the same time greater protection of their right to privacy.“*

The argument of critics, that access to the data of all EU citizens by hundreds of thousands of European security personnel is in itself a security risk, is turned on its head in the most absurd way. The repression authorities are no longer confronted with the problem of gathering data in light of the many ways in which data is currently generated and collected: registration and revenue offices, provider data, banks, user profiles on the internet (MySpace, Facebook, Second Life), e-government, travel profiles, telecommunication surveillance, video surveillance, GPS. To be ‘armed’ in “the era of cyber-space” rests on adequate administration if data and the ability to make ‘sense’ of it. The talk is of an “almost limitless amounts of potentially useful information”.

> *“Information is the key to protecting the public and in an increasingly connected world in which public security organisations will have access to almost limitless amounts of potentially useful information. This is a challenge as well as an opportunity – public security organisations will need to transform the way they work if they are to master this data tsunami and turn it into intelligence that produces safe, open and resilient communities.“*

To facilitate European-wide data exchange interior ministers have already agreed to 6 of 49 “types of relevant information”. DNA, finger-printing, Ballistics, vehicle registration, telephone numbers and registration data. This catalogue is to expanded to a “Top-Ten” list in 2009. A severe problem are the different standards that exist within the member states with respect to hardware, software, format, but also the systematisation of data. The ‘Future Group’ would like to set up a “European Union Information Management Strategy” (EU IMS) in order to develop standards and promote system cooperation.

> *“The key to effectiveness will be using technology to connect the capabilities of a multitude of stakeholders and ensure the right information gets to the right person in the form they are best able to use.“*

An “interoperable platform” should facilitate better communication between the different police forces of the member states and the EU institutions. Existing databases such as the Schengen Information System II (SIS II), the Frontex-Portal Border Technet, the Europol Network European Information System (EIS) or the biometric visa database VIS are to be connected as “convergent networks”. With this, a surveillance network of previously unknown dimensions will be created, and will function as a central nodal point of Europol, as a “competence centre for technical and coordinative support“. In the long-term, Europol is to develop a “security partnership with Interpol” (the second-largest international organisation, second only to the UN) and cooperate closely with SitCen, the exchange platform for intelligence.

The paper proposes that data exchange should also be extended to “third countries”. The focus for this is the USA whose regulations already allow the release of data to other authorities and countries. The intention is to decide upon setting up “Euro-Atlantic area of cooperation” by 2014. Real-time data flows will play an important role in this, whereby access to large volumes of data by the authorities will be possible from any place in Europe, also during police operations. For this,

broadband mobile networks are being built. At the same time, technologies such as RFID, WLAN or Bluetooth will enable live-protocolling of behavioural patterns.

> *"Special investigative techniques should be placed higher on the agenda. [...] Member States should prioritise investment in innovative technologies that enable automated data analysis and improve real-time collaboration. Research in these areas should be encouraged, ensuring that ideas can move quickly from a research context to practical implementation."*

Automated data analysis functions in the following way. Computers systematise data of persons, objects or criminal offences (as processes that run in the background). They can be constructed as relationship diagrams. These are compared and risk analyses constructed from them. The software can also process audio files such as those obtained through telecommunication surveillance or excerpts from interrogations. The result is a visualised "mapping" of complex relational structures. Three-dimensional pictures are created through the generation of different information 'layers' that are placed on top of each to identify 'clusters'. This software can provide assistance in decision-making that corroborates previous data, incorporating simulations (e.g. larger police operations at summits or sport events). Used in real-time these can identify "suspicious telephone conversations"or, combined with biometrics, can identify unusual behaviour or clothing features.

Such risk analyses mark a shift towards a "proactive approach" to policing. A population or certain groups can be placed under general suspicion and researched by machines. With this, police and intelligence services intend to foresee crimes. Here, a fundamental paradigm shift within police operations is taking place. To date, police can only take action if a crime has occurred or if there is evidence of a criminal offence; thus, legislation regulating police operations will have to be changed to accommodate these shifts towards preemptive strategies.

The strategy paper of the Future Group advises the EU to use "preventive and repressive" measures against "terrorist threats", complimenting these with "proactive" measures with the help of civil society and business. A particular focus is the internet. Beside surveillance centres, active intervention through the internet should also contribute to what is being called "de-radicalisation" by making use of "cultural intelligence", taking account of "cyber-language". Yet this is not the whole of the information war of the cyber-speaking interior ministers. They also advise on the need for media strategies,

> *"focusing on inter-cultural dialogue and developing a clear and convincing positive message to different communities in Europe and abroad – possibly even in non-European languages, with regard to European core values of good governance, fundamental rights and safeguarding of peace and freedom."*

Data protection is strongly under-represented in the strategy paper. A general justification for new measures is provided by the supposed desire for people to want more surveillance and control of their lives:

> *"Ensuring greater public understanding of the benefits of data sharing between Member States should be a priority. The strategy should include a commitment to make clear to European Union citizens how information will be processed and protected, on the basis of proportionality and necessity."*

This new five year plan for Home Affairs is accompanied by a concerned look to the political developments within the EU: In Spring, a new Commission president will be decided upon, in June a new parliament will be elected. The publication of the strategy paper, 'Freedom, Security, Privacy – European Affairs in an Open World' intends to assist the new (presumably more right-wing) parliament in ratifying the changes in EU security policy without much resistance. Their conclusion, that the strategy paper should only be understood as "reflections and ideas" should be understood as a mere euphemism in light of the concrete nature of the plans proposed in the paper.

A Comprehensive Approach for Social Movements

Proposal for a Campaign against the EU

The phenomena described in the assessments of current proposals and plans for a transformation of EU and NATO approaches in the area of security policy present social movements with great challenges. Clearly, there is a logic at work here that is engulfing ever-more of social life. Resisting the double challenge of militarisation and securitisation cannot be left only to peace movements, antimilitarist, civil liberties or anti-repression groups – particularly when many of these operate on a national level. In a society where economic processes divide people into winners and losers, where

in a global context disparities are intensifying, plans for a more ‘Comprehensive Approach’ and ‘Homeland Security’ are attempts to create order and security for capital accumulation through military means. Where conflicts arise that cannot be solved through integration, securitisation provides the solution. The effects of this are experienced in the every-day as increased social control, the criminalisation of poverty and the fight against migration. Under a securitisation logic, every social conflict is seen as a potential threat. A politics of fear expressed through images of enemies and threat scenarios legitimises authoritarian and militarised strategies of domination. More control, exploitation and a state of exception become the norm.

To highlight and counter these developments, we propose a “Summer of Resistance 2.0”, a ‘comprehensive approach’ of social movement mobilisations against the NATO summit in April 2009, the Swedish EU presidency in the second half of 2009 and the G8 summit in Italy. A number of European groups are engaging critically with EU politics and are affected by the consequences of them. Themes, and thus participants to such a campaign could be: Migration, Filesharing networks, Alternative providers, Terrorism Cases, Data retention, Critical lawyers, Civil liberties groups, Lissabon Treaty, Antimilitarism, Peace movement, NATO, G8, Economic Partnership Agreements (EPAs), Climate, Environment.

Networked and collectively acting social movements can put a limit to these security phantasies and strategies in order to push open the space to make alternatives visible. We would like to discuss this proposal at future encounters of social movements. Please send us any feedback at mail@gipfelsoli.org.

1 General a.D. Klaus Naumann (Germany), General John Shalikashvili (USA), Field Marshall Lord Peter Inge (UK), Admiral Jacques Lanxade (France), General Henk van den Breemen (NL).

2 All citations in this section are from the Future Group paper unless otherwise stated.

3 See http://europa.eu/scadplus/leg/en/lvb/l14579.htm

Background

- European Home Affairs in an Open World: http://euro-police.noblogs.org/gallery/3874/eu-futures-jha-report.pdf
- “Towards a Grand Strategy for an Uncertain World”: http://euro-police.noblogs.org/gallery/3874/grand_strategy.pdf
- Analysis of Statewatch, “The Shape of Things to Come”: http://www.statewatch.org/analyses/the-shape/of-things-to-come.pdf

Gipfelsoli September 2008

http://gipfelsoli.org | http://euro-police.noblogs.org

25th Chaos Communication Congress
Nothing to hide
Berlin, 2008-12-30

La Quadrature du Net

Campaigning on Telecoms Package

Pan-european activism for patching a pirated law

by
Jeremie Zimmermann
<jz@laquadrature.net>

Whatever you do will be insigni cant, but it is very important that you do it.
Mohandas Karamchand Gandhi

Contents

Prelude

It is clear for every participant to 25C3 how digital technologies and the Internet are crucial and structuring for the future of our societies. Yet, as the technology and the network are neutral per se, we all know they can be used in ways that range from the most fantastic tool ever invented for knowledge sharing and human enlightenment, to the most frightening environment of surveillance and control. It is our responsibility, to us people of the network, who grew along with and built these technologies, to allow the first to arise and block the latter from becoming real. We may even be the only ones able to do so.

In the last ten years, corporations whose obsolete business models are based upon control have increasingly used their power to strengthen the vision of a technological environment being used against its users, their freedoms and their privacy. They use lies, ignorance and corruption to achieve their goals. Political influence and legislative tricks, mastered by the powerful lobbies representing these industries, are their primary way of turning their dreams, our nightmare, into reality.

When arose one of the most obscurantist and stupid law we ever saw in our activist experience, the French Olivennes bill, HADOPI, or graduated response, when we saw the manipulations of a few to propagate it to the whole Europe, we knew it was time to react.

It was time to combine our knowledge of these processes, our will of freedom, and the use of the network. It was time for documenting all the wrongdoings of our powerful opponents. It was time for working collectively towards the primary objective of raising awareness on fundamental rights and freedom in the digital environment. It was time for La Quadrature du Net.

La Quadrature is a toolbox for everyone to understand the political processes our enemies pervert in order to gain more control. It is also designed, based on our collective experience in campaigning, to allow everybody to take action through powerful campaigning tools. The formidable cultural and humane diversity of a pan-European network of friends and siblings made it more powerful than we ever imagined.

We believe in the equalitarian design of the digital technologies we build and use, we know that knowledge and sharing are the keys for building a better society.

This is the story of our first battles, our first victories. We expect it could give hope to the ones thinking that there is nothing we can do to change anything.

It is very important to remember that this is a long-term fight. A fight for freedom and equality. A fight we know we can only, altogether, win.

Jérémie Zimmermann

"Whatever you do will be insignificant, but it is very important that you do it."
Mohandas Karamchand Gandhi

Who are we ?

Who is *La Quadrature du Net*

La Quadrature du Net is a group of citizens providing informations about legislative projects that threaten civil liberties as well as economic and social development in the digital age.

La Quadrature du Net informs citizens, public authorities, organizations, corporations.

It works with everyone to elaborate balanced alternative solutions.

La Quadrature du Net is supported by French, european and international NGOs including the Electronic Frontier Foundation, the Open Society Institute and Privacy International.

Why this name?

We believe that the promoters of the projects we are opposing are trying to solve a problem similar to squaring the circle. They do not understand that we changed era, that some approaches are outdated, that we must collectively rethink our approach to the control of information.

We believe that it is impossible to effectively control the flow of information in the digital age by the law and the technology without harming public freedoms, and damaging economic and social development. This is what we call squaring the net (in french: *La Quadrature du Net*).

We agree with the idea expressed by Jacques Attali related to the draft of the Olivenne law, when he presented the Committee's report about economic growth to the parliament: "*we cannot develop economical growth by installing surveillance and tracing.*"

For the records, according to the online encyclopedia Wikipedia, squaring the circle is a classic problem of mathematics appearing in geometry. It is one of three major problems of the ancient world, with the trisection of the angle and the duplication of the cube.

In the oldest mathematical text found, Rhind papyrus ($\approx$ 1650 B.C.), The scribe Queen Ahmose already proposed an approximate solution of the problem. However, we had to wait until 1882 for the German mathematician Ferdinand von Lindemann to show the transcendence of π, thus proving that it's impossible to square the circle: it is impossible to construct, using only a ruler and a compass, a square whose surface is exactly equal to the surface of a disk.

The question now is how many centuries will it take for the law makers to return to reason? Are we in the field of law and ICT closer to BC -1650 or 1882?

They support *La Quadrature du Net*

Electronic Frontier Foundation

Open Society Institute

Privacy International

April
(France)

Asociacion de Internautas
(Spain)

Big Brother Awards
France

Com a CC OO
(Spain)

Creative Commons
France

French Data Network
(France)

Globenet
(France)

IT-Political Association of Denmark

MarsNet
(France)

Musique Libre
(France)

Netzwerk Freies Wissen
(Germany)

Open Rights Group
(United Kingdom)

Qwartz Electronic
Music Awards
(France)

support us ! write to contact@laquadrature.net

Campaign on Telecoms Package and graduated response

Squaring the net - March 24, 2008

The French parliament will soon debate a draft law "*about the High Authority for the dissemination of works and protection of the rights on the internet*". The project incorporates the recommendations made by the former CEO of FNAC (biggest CD and DVD retail stores in France), Denis Olivenne. One of the key measures is to sanction a presumed violation of copyright by cutting access to the Internet.

This sanction would be imposed by a so-called "independent" administrative authority. Reported by private actors monitoring the Internet, the alleged offencers would then be remotely recognized by administrative officers who have access to personal connection data. All this without supervision of a judicial authority. So far, only law enforcement agencies fighting terrorism have such exceptional power, on a temporary basis, until the end of 2008.

This staggering bill, prepared under unbelievable conditions - Was a mission on GMOs ever given to Monsanto? - is indicative of a dangerous headlong flight for democracy, society and economy.

In France, many laws and regulations about digital technologies have been adopted in recent years, without prior consultation or assessment of the existing texts, under lobbies pressure. The Olivenne draft law is being drafted while the implementation report of the very controversial 2006 law on copyright, due in February 2008, has not even begun.

The previously adopted texts are out of step with the reality of usage and techniques. Just voted, they are obsolete, unenforceable, ineffective. Lawyers don't even understand them. It's squaring the Net: the devil does not want to return into the box! And what if the devil was the box itself?

Asking the question is by itself heretical. Just appointed, Denis Olivenne announced that it was out of the question to legalize sharing of music and films for a fee distributed among artists. Such a blanket license was provisionnally voted in late 2005 by MPs from all sides, before the government, pressed by the publishers, brought the Parliamentary majority back in line to withdraw it.

The Olivenne project therefore drives the government into a dogmatic, almost obscurantist and authoritarian drift. It takes inspiration from obsolete oracles, considering the progress as a threat rather than an opportunity. The proposed step is however an unprecedented step backwards.

The Swedish ministers of culture and justice, who recently rejected a similar proposal, made no mistake, stating that "*Many have noted that shutting down an Internet subscription is a wide-reaching measure that could have serious impacts in a society where access to the Internet is an imperative welfare-issue.*" Especially in the case of "*triple play*" offers, where the telephone and television could be cut too. The electronic social death of entire families on behalf of copyright? Beaumarchais, Victor Hugo, Jean Zay must be turning in their graves.

And what about the extension of emergency measures intented for combating terrorism, so that private companies can hunt the Internet users and circumvent the judicial authority? Who can believe that an Internet user punished in this way will turn to spend money in FNAC stores? Especially when one considers that disconnected users will still have to pay for their connection!

As for the cost to the taxpayer and the economy, it remains unknown today. No economic impact analysis has been carried out. To monitor, threaten and repress millions of people via a parallel justice has a price, however. To disconnect homes and businesses, as well. Public finances and all subscribers will have to bear the cost of this surveillance.

There are, however, other solutions: to legally secure democratic and creative uses of the Internet, enabling web entrepreneurs to innovate safely, to review the existing mechanisms for wealth distribution

and finally admit that the punitive approach and dialogue reduced to (only between?) a few lobbies lead to market authoritarianism.

It is impossible to effectively control the flow of information in the digital age by the law and technology without causing serious harm to civil liberties and hindering economic and social development.

It took 3000 years to demonstrate that it was impossible to fit circles into squares, with a ruler and a compass, without loss, because of the transcendence of π. As squaring the circle in its time, squaring the net will be only overtaken by changing tools and perspectives. Will it take 3000 years for our elites to understand it?

The debate is open - March 27, 2008

While the Olivenne bill[1] aimed at building the "graduated response" must be voted before the summer in France, a group of citizens launched "*La Quadrature du Net*" to alert on this and other equally disturbing governmental projects, and to make alternative proposals.

Since the beginning of the year, the french government has announced various projects relating to Internet:

- Olivenne bill aimed at building "graduated response"[2] allowing an administrative authority to cut down internet access that has been repeatedly used for downloading music or video without authorization;
- The proposed extension of the powers of the CSA (french media authority) to Internet.
- Draft proposal for an administrative marking of websites.
- Blurry plan to combat cybercrime.
- Draft decree extending the retention of connection data.

The government has indicated that it wanted several of them, including the Olivenne bill, to be adopted before the French presidency of the European Union (July 1st), and that these measures are widespread at the euroepan level in the wake.

"*Widespread monitoring of Internet, including by private actors, mandatory over-referencing of 'accredited sites' by the search engines, administrative supervision of content hosters and publishers, content filtering and internet access disconnection without trial ... These projects draw a declining democracy, a political control of the internet, a Big Brother society. By no means a model for Europe.*" Said Christophe Espern, co-founder of the initiative.

Emergency measures planned to fight terrorism could be extended over time and to other areas, in an effort to preserve outdated models, without a real democratic, open and transverse debate. Such a debate is essential regarding the ethical, social and economic issues caused by the "*digital revolution*".

Citizens who believe that there are other answers to this challenge therefore decided to launch an initiative to inform the public and the government, and open the debate. They chose the name *La Quadrature du Net* because they believe that, for years, the lawmaker is trying to solve a problem similar to squaring the circle.

The site http://laquadrature.net relays information about the projects and future actions of the initiative. Synthetic information kits are online, they will be updated throughout time. Other analysis, op-eds and proposals will follow. Users can stay informed by subscribing to the mailing list and participating more directly to the initiative[3].

[1] The Olivenne bill is named after Denis Olivenne, CEO of FNAC (one of the biggest CD, DVD and online music retailer in France), commissioned by Nicolas Sarkozy to find solutions to fight unauthorized downloading.

[2] EFF about the "*three strikes*": http://www.eff.org/deeplinks/2008/03/three-strikes-three-countries

[3] http://www.laquadrature.net/wiki/How_to_Help

European Parliament rejects graduated response - April 10, 2008

The European Parliament adopted a resolution this morning which commits the member states - therefore France - "*to avoid adopting measures conflicting with civil liberties and human rights and with the principles of proportionality, effectiveness and dissuasiveness, such as the interruption of Internet access.*" This vote proves that the system of graduated response that Nicolas Sarkozy wants France to adopt quickly and to extend to Europe during the French Presidency of the EU, is seen as contrary to human rights by a majority of MEPs.

This vote is a strong signal sent towards the French governement. It comes in support of the position of the Swedish government which had already rejected the graduated response. The rapporteur Guy Bono brought this resolution, which is supported by members from all political sides, stated yesterday in plenary:

"*On this subject, I am firmly opposed to the position of some Member States, whose repressive measures are dictated by industries that have been unable to change their business model to face necessities imposed by the information society. The cut of Internet access is a disproportionate measure regarding the objectives. It is a sanction with powerful effects, which could have profound repercussions in a society where access to the Internet is an imperative right for social inclusion.*"

La Quadrature du Net, that wrote a 3 pages letter monday to MEPs, welcomes this vote. We thank all the elected officials who voted for the amendments that led to this result. We also thank all freedom defense organizations we worked with to raise awareness among MEPs on the topic, through e-mail and telephone. We invite the French Prime Minister, Francois Fillon, to take this vote into account, and therefore not to submit the Olivenne Bill (french graduated response) to the French Parliament.

As explained in the report by the professor of criminal law Jean Cedras that the minister Renaud Donnedieu De Vabres had sought to bury in his time, "*the idea of an automatic graduated response, as tempting as it seemed [to the French government and the right-holders], should be abandoned.*".

Will France Introduce the Digital Guillotine in Europe? - April 23, 2008

Ever since DADVSI, the French implementation of the European Union Copyright Directive (EUCD), Internet users in France have faced increasingly disproportionate threats of punishment for claims of copyright infringement. The latest scheme promoted by the content industry against unauthorized sharing of music and films on Internet is called "*flexible response*" or "*three strikes, you're dead*".

The principle is simple: repeated infringements result in getting cut off from the Internet. A claimed copyright infringer is identified by automated Internet traffic filtering and by a rightsholders' denunciation. After a complaint to the ISP, a letter is sent warning the alleged infringer that he might be being cut off from the Internet. In the early versions of this scheme the punishing fines were to be sent out automatically, but the fines were later replaced by the proposal to cut off Internet access instead.

The proposal further includes the creation of an administrative authority responsible for enforcement, making sure the disconnected Internet users are not able to use the Internet again for a set period of time. The scheme is unclear as to the possibility to appeal a mistaken claim or to whether a punished Internet user can also be sued in a civil law suit. State authorised software for securing Internet connections has been proposed as one solution to uphold legal protection of innocent citizens.

Promoted by the French President Sarkozy, flexible response has become known as the suggested main measure in the Olivenne report, which until recently was generally thought to provide the basis for both French and European legislation to come. But not any more; a majority vote in the European Parliament on the 10 April 2008 suggests otherwise.

In contrast to the French solution, the Swedish government has rejected such a regime as disproportionate. The Swedish Ministers of Justice and Culture concluded in March 2008 that ostracism from the Internet as punishment in a society whose daily activities are increasingly intertwined with the digitally networked environment is not proportional to the infringement of copyright, especially without intention for commercial gain. Its only justification, that of deterrence, has been shown repeatedly to be ineffective. The Swedish Government has also pointed out that large owners of media content should "*not use the copyright laws to defend old business models*" but should do more to provide attractive alternatives to unlicensed filesharing services.

In an effort to bring the Swedish Government's policy decision to the European discussion, and to oppose Sarkozy's plans, MEP Christofer Fjellner (EPP) initiated a cross partisan platform together with former French Prime Minister MEP Michel Rocard (PSE) and MEP Guy Bono (PSE), Rapporteur for the European Parliament's Report on Creative Industries. Together they signed, with more than 90 MEPs supporting them, an amendment to the report which effectively rejects flexible response:

"*Calls on the Commission and the Member States to recognise that the Internet is a vast platform for cultural expression, access to knowledge, and democratic participation in European creativity, bringing generations together through the information society; calls on the Commission and the Member States, therefore, to avoid adopting measures conflicting with civil liberties and human rights and with the principles of proportionality, effectiveness and dissuasiveness, such as the interruption of Internet access.*"

Following the tabling, the internal debate over the amendment intensified and was complicated when the liberal group (ALDE) unexpectedly asked to split the amendment into two parts and vote on them separately, presumably to save flexible response. The first part related the Internet and the importance of rights and proportionality had overwhelming support, while the second part explicitly mentioning "*interruption of Internet access*" was harder to support for many French MEPs who would not go against the explicit will of their own government. However, with a close vote on the second part, the amendment ended up being passed in its entirety.

In the aftermath of the European Parliament's decision, the French Minister of Culture, Mrs. Christine Albanel, has clearly announced her intention to move on with the Olivenne proposal. She is currently planning to put it to vote in the French Parliament before the summer. Accordingly, to this date, there are no indications that Sarkozy's decision to use the French Presidency to propagate this scheme at the European level has been revised.

In Brussels it is also unclear whether the initiative by MEPs Fjellner, Rocard and Bono will have an impact. Neither of them is listed as a speaker or moderator at the next High Level Conference on Counterfeiting and Piracy on the 13 May 2008, while MEPs who voted in favour of flexible response are. It is a matter of public interest to ensure that there is a balanced debate and that seats are reserved for politicians representing the European Parliament's position on what otherwise risk to be a very controversial conference.

In media and the political blogosphere the impact of the vote is increasing. Of particular interest is the correlation between the Member States with a well developed Internet infrastructure and the way their MEPs voted: the digitally advanced Nordic countries have all clearly rejected the French digital guillotine.

Amendment 1 by Christofer Fjellner and amendment 2 by Michel Rocard and Guy Bono and others. Report Cultural industries in Europe. Rapporteur Guy Bono (2.04.2008): http://www.europarl.europa.eu/sce/data/amend_motions_texts/doc/P6_AMA(2008)0063(001-001)_EN.doc

Privacy: Film industry pirates European law - May 14, 2008

***La Quadrature du Net* is worried about amendments endangering privacy tabled by the rapporteurs of the Culture Committee of the European Parliament. They fit into the consideration of two proposals framework directives known as "*Telecoms Package*"**[4]

These amendments are aimed at injecting in the Privacy and electronic communications (E-Privacy) directive a device known as the flexible response or three strikes. It is allowing producers to engage in police duties, and access providers to punish the public without going through the judicial authority.

As evidenced by a note of the french film lobby circulating in the European Parliament, several of these amendments are measures drafted directly by the film industry that seeks to break a recent ruling by the European Court of Justice. They significantly reduce the level of protection of privacy and personal data in Europe.

La Quadrature du Net is concerned about other amendments intended to legalize cultural industries' spyware, to institutionalize their influence, or to enable them to determine what wireless technologies

[4]- The Telecoms Package: http://www.laquadrature.net/wiki/Telecoms_Package
- Guardans amendements: http://www.europarl.europa.eu/meetdocs/2004_2009/documents/pa/718/718636/718636en.pdf
- Mavrommatis amendements: http://www.europarl.europa.eu/meetdocs/2004_2009/documents/pa/718/718261/718261en.pdf

will be used by the public. The last amendment studied denies the existence of a public's right to redistribute the public domain or free licensed works.

"*These amendments confirm once again that the 'flexible response' goes against the fundamental rights and the principle of proportionality. How MEPs could have tabled such dangerous legislative riders? The reform package Telecoms aims to strengthen the protection of users, not weaken it! The film industry wants to pirate European law!*" Says Christophe Espern, spokesman for *La Quadrature du Net*.

La Quadrature du Net therefore invites MEPs to confirm their 10 April vote condemning the graduated response, by rejecting these amendments.

Analysis by *La Quadrature du Net*:
http://www.laquadrature.net/files/note-CULT-quadrature-13052008-en.pdf

MEPs want to torpedo the Free Internet on July 7th - July 1, 2008

One week before a key vote in the reform of European law on electronic communications ("*Telecoms Package*"), *La Quadrature du Net* denounces a series of amendments aimed at closing the open architecture of the Internet for more control and surveillance of users.

European Internet users could be blocked from lawful activities by a mandatory spyware, in the interests of their security. The right to use free software for internet access would therefore not be assured anymore. The neutrality of the Internet is also directly attacked, as is the principle that technical intermediaries have no obligation to prior surveillance of contents. Other amendments will de facto enable administrative authorities to obligate ISPs to work with content producers and rights-holders' private police, including the sending of intimidating messages, with no judicial or regulatory oversight.

This measure goes further than the French "*graduated response*" project, which has been subject to widespread opposition, including by the European Parliament on April 10th. That is undoubtedly why those amendments have turned up on early july, and why those draftings them use subtle rhetoric and crossed-references to make the overall text harder to understand (more than 800 amendements on 5 directives were tabled).

"*The politicians who engage in these summer manoeuvres dishonour Europe and their mandate. They rely on the fact that nobody watches them a week before Parliamentary holiday, to divert the Telecoms package from its primary objectives of consumer protection. They pave the way for the monitoring and filtering of the Internet by private companies, exceptional courts and orwellian technical measures. It is inconceivable for freedom but also for European economic development. We call on all MEPs to oppose what they have already rejected.*" said Christophe Espern, co-founder of *La Quadrature du Net*.

These torpedo amendments are currently subject of a series of secret, back-room negotiations between a handful of MEPs who do not always understand all the implications of these issues. Accomplices of lobbyists who hold the pen are in every political party. Instructions for the plenary vote will be established this week for a vote in IMCO and ITRE committee on Monday, July 7th.

At this stage, citizens must act urgently and en masse, to make their MEPs understand, a year before the elections, the possible consequences of their actions.

More infos on how to act and get mobilized: http://www.laquadrature.net/wiki/Mobilisation_Paquet-Telecom

The "*Telecoms Package*": out of the shadows, into the light - July 10, 2008

On Monday, July 7th, the IMCO and ITRE committees of the European Parliament passed the review of European telecommunications law known as "*Telecoms Package*". All of the amendments damaging to the Internet, that were condemned by la Quadrature du Net and numerous organizations, have been voted through.

La Quadrature du Net continues to maintain that these amendments aim to close the open architecture of the Internet, with the result that there will be increased surveillance of users. They open the way to the regulation of users via the Internet Service Providers (ISPs) under the control of national

regulators, instead of a judiciary authority. They directly attack net neutrality by allowing these operators to constrain legal activities of users (such as the use of free software), in order to promote their own services or those of companies dominant in the market for content or software.

However, *La Quadrature du Net* is delighted to see the strong citizen mobilization which emerged in parallel with this committee vote. Thousands of European citizens have written to MEPs. The European press has actively covered the vote. Numerous MEPs became aware of the stakes and the threats to infringe on basic rights, the open architecture of the Internet, Free Software, competitivity and European informational sovereignty.

During the session, many MEPs have underlined the importance of citizen mobilization and how it helped to highlight many problematic elements of the Package, such as privacy. Others have complained about getting the messages, undoubtedly upset to see that their proposals did not meet the publics expectations.

Some of the European Parliament rapporteurs on the Telecoms Package have also re-read their proposals and admitted that, when viewed in the overall context, they were problematic and should be reworked before the final plenary vote .

We are pleased with the impact our campaign has made, but also disappointed that IMCO committee MEPs have adopted amendments they know to be damaging or that they could not study at length.

How else can one explain that shadow rapporteurs tried to orally amend, as an emergency measure, compromise amendments they had accepted shortly before? And how a member of the Parliamentary Committee could take note of compromise amendments that were unknown before he entered the session?

But we are pleased to see that today that several MEPs, including rapporteurs of the IMCO Committee, request a postponement of the Telecoms Package vote in the plenary, in order to have more time to examine it. They particularly wish to delete certain amendments adopted in LIBE Committee that *La Quadrature du Net* has denounced as particularly threatening to privacy. Others want to reintroduce judicial authority into the measures, which is essential, as the effective protection of the net neutrality.

"*The dialogue between users and MEPs is really gratifying, as awareness on these issues that takes place in many parliamentary groups. It is, of course difficult for an elected official to acknowledge that the amendments he has tabled, intertwined with those of colleagues from other groups are catastrophic. But I can not believe that once informed, the European Parliament does not correct itself.*" said Christophe Espern, co-founder of *La Quadrature du Net*.

La Quadrature du Net therefore calls upon the citizens to continue to contact MEPs to invite them to ask their parliamentary group presidents to postpone the vote in plenary, normally scheduled on September 2, to October. MEPs must have the time to study this complex and transversal dossier, and so must citizens.

Five directives are modified by three framework directives and more than a thousand amendments have been tabled so far. Many weeks, excluding Parliamentary holidays, are needed to identify all the problems and prepare amendments that effectively protect the rights of citizens and the open architecture of the Internet.

La Quadrature du Net will soon publish a detailed analysis of the finally adopted amendment.

Telecoms Package: vote postponed - July 11, 2008

Planned for September, 2nd, the vote of the Telecoms Package was postponed by the Conference of Presidents (which is made up of the chairs of the political groups and the President of the European Parliament). The debate will take place in Strasbourg, September, 2nd but the vote would be planned during the session which will start September, 22nd. For several MEPs, this is an unusual situation. More infos soon.

Telecoms Package: the spectre of the graduated response hangs over Europe - September 3, 2008

MEPs, representatives of the European Commission and Council have discussed yesterday[5] in plenary session, in Brussels, the reform of European law on electronic communications (Telecoms Package).

Among the fifty Members of European Parliament (MEPs) who spoke, only a minority defended the provisions relating to copyright and content that were introduced by some parliamentary committees last July. Several MEPs conversely requested that these provisions be removed as hazardous to human rights.

The very same morning, the European Data Protection Supervisor[6] (EDPS) published a critical comment[7] on several of these amendments (amendments 9, 30, 76, 81, 112, 130 and 134 brought in IMCO by British Conservative MEP Malcolm Harbour).

The EDPS, as an independent EU authority in charge of personal data protection, considers that these amendments are an open invitation to a "*mass surveillance of Internet users*" and to "*lay the foundations*" of the graduated response. The EDPS confirms by the way the analysis that *La Quadrature du Net*[8] had published in July, like many other notes that were sent to MEPs.

La Quadrature du Net therefore strongly denounces the remarks made by the appalling representatives of the French presidency, Luc Chatel and Eric Besson, who explained that the graduated response respects the rights of citizens. How can they keep holding this speech in view of the EDPS advice?

"*The EDPS report is a real blow to promoters of the graduated response; it is incredible that some MEPs and the French presidency still act as though nothing happened. In any case, Parliament has the opportunity to show that the citizen-protective Europe is not a fiction.*" said Christophe Espern, co-founder of *La Quadrature du Net*.

Jérémie Zimmermann, also co-founder, adds: "*MEPs must recognize that the purpose of the Telecoms Package is to protect consumers, rather than giving away their privacy. Citizens should call their representatives in the European parliament and explain them.*"

La Quadrature du Net therefore invites European citizens to contact their MEPs so that they vote, during the Telecoms Package vote on September 23rd, for the removal of the amendments the EDPS criticized. MEPs must protect privacy, whatever the French presidency and the lobbies it serves require.

Telecoms Package: protect the free and just society! - September 19, 2008

The crucial first reading vote on the "*Telecoms Package*" will take place in the European Parliament, in Brussels, on Wednesday, Sept. 24th. Even if some noticeable progress was made, some dispositions of these internet regulation directives still pose an important threat to civil liberties and fundamental rights[9]. *La Quadrature du Net* calls for its supporters to mobilize on the amendment 138 tabled on the Trautmann report[10] to guarantee that "*graduated response*" could not emerge in Europe.

"*We want Europe to protect citizens, as stated in the primary objectives of the Telecoms Package. Conversely, these directives must not on the opposite erode individual rights and liberties.*" declares Jérémie Zimmermann, co-founder of *La Quadrature du Net*.

Unfortunately, even if some amendments were positively reworked and neutralized according to the recommandations of the European Personnal Data Supervisor (EDPS), rapporteurs of the texts most of the time didn't follow the primary recommendations. Instead of deleting the concerned problematic

[5] http://www.europarl.europa.eu/sce/server/internet/cre/sce_cre_02.jsp

[6] http://www.edps.europa.eu/EDPSWEB/edps/lang/fr/pid/1

[7] http://www.edps.europa.eu/EDPSWEB/webdav/site/mySite/shared/Documents/Consultation/Comments/2008/08-09-02_Comments_ePrivacy_EN.pdf

[8] http://www.laquadrature.net/wiki/AnalyseCompromisVersionDiteFinale

[9] http://www.laquadrature.net/wiki/Telecoms_Package_Plenary_Amendments

[10] http://www.europarl.europa.eu/sce/data/amend_motions_texts/doc/P6_AMA%282008%290321%28138-139%29_EN.doc

parts of the text, they chose to try to rewrite them. The result is a vague and broad text[11], introducing new concepts in the European law (such as the notion of "*lawful content*").

The main risk is that this fuzzy wording be used by some Member States to authorize administrative authorities to restrict freedom of expression and information of internet users suspected of unlawful file sharing, without any prior judicial ruling.

This is not a fantasied, illusory risk. For instance, the French parliament shall consider as soon as november a draft law transfering repressive power to an administrative authority acting upon request from the cultural industries' representatives.

It is essential to us that the European Parliament eliminates that risk that could potentially jeopardize the proportionality principle as well as the separation of the powers, but also weaken the acceptability of necessary criminal enforcement measures. Measures that national public authorities can implement to fight terrorism or child pornography shall not be extended to issues related to non-profit music or movie sharing over the internet between individuals.

La Quadrature calls all its supports, consumers, citizens to urgently help raise attention among Members of European Parliament (MEPs)[12] on those issues, and ask them to vote for the Bono/Cohn-Bendit/Roithova amendement (138) in order to guarantee that civil liberties will remain protected by Justice.

Telecoms Package: European democracy's victory already threatened - September 26, 2008

***La Quadrature du Net* welcomes the adoption, in the first reading, of several amendments correcting major problems in the Telecoms Package, as well as the rejection of the most dangerous amendments.**

Members of the European Parliament have shown today their commitment to privacy, the protection of personal data, and principles of proportionality and separation of powers.

La Quadrature thanks all MEPs who have worked in this direction, and all citizens who mobilized en masse to alert their delegates on these issues. We'd like to thank particularly the MEPs who have been able to reconsider their positions as they became aware of the risks to the rights and freedoms of their fellow-citizens.

However, La Quadrature calls for watchfulness.

Some vague formulations remain in the Telecoms Package. They do not prevent transposition into national laws affecting network neutrality. For example, the concept of "*lawful content*" is unknown in European law; its definition is left to the Member States. It must be completely removed, in order to complete the cleaning already made on this Package.

In addition, European Commissioner Viviane Reding has already announced she would require, on behalf of the European Commission, the withdrawal of Amendment 138 voted this morning.

Amendment 138 yet reaffirms a fundamental principle which should bring together all Europeans attached to the pillars of Europe, since it states that no restriction on the rights and freedoms of end users can be taken without prior decision of the judicial authority - safe when public safety is concerned (prevention of harm to persons).

La Quadrature therefore denounces a completely unsuitable request from Mrs. Reding, under the basic democratic principle recalled in the amendment (i.e. the separation of powers), but also under the parliamentary plebiscite it collected (574 MEPs for, 73 against).

This is indeed a self-evident democratic issue that a technocrat cannot deny, whatever her will to serve the liberticide plans of the French government and the entertainment lobbies. It is indeed obvious that she is trying to save the French "*Graduated Response*" project (also called "*three strikes and you're*

[11] For an example, look at amendment 191 (http://www.laquadrature.net/wiki/Telecoms_Package_Plenary_Amendments#Amendment_191_--_replace_Amendment_9) introducing recital 12(c) about "*cooperation*" between administrative authorities, internet service providers, and cultural industries' representatives. Nowhere is clearly closed the door to target messages based on surveillance of individual behaviour of internet users. This is a major part of the "*graduated response*" scheme that isn't precisely framed here.

[12] http://www.laquadrature.net/wiki/Political_Memory

out"), at which amendment 138 directly aims.

This is not the first time, however, that MEPs have stressed the illegality of the so-called three-strikes approach under community law. Last April they have already passed a resolution on cultural industries, which condemns the French project because of its disproportion.

This position was also expressed by the European Data Protection Supervisor early September. The French government should therefore be the one to reconsider its position, instead of asking the Parliament to give up on its duty to protect the fundamental rights of European citizens.

La Quadrature now asks all EU citizens to write to the president of the European Commission and to their government, and to make sure that the amendment 138 will be preserved by the Council. They also point out for lobbies and the French Minister for Culture that it is ridiculous to pretend MEPs did not target the graduated response.

For months, France has interfered with the Telecoms Package by trying to introduce copyright-related amendments in it. Fortunately, European citizens massively wrote to or called their MEPs to ask them to preserve their fundamental rights against this offensive. As a consequence, MEPs are totally aware of the stakes.

For instance, La Quadrature sent this note to all MEPs the day before the vote, to explain exactly why they consider the amendment 138 essential:

Letter sent to all MEPs on september 23rd

Dear MEP,

On September 24th the European Parliament will be examining the first reading of the bill proposing the reform of the law on electronic communications, known as the Telecoms package.

At the beginning of this summer this bill aroused a great deal of controversy. Several amendments which were adopted during commissions were denounced by some NGOs, as they would lower the level of data protection in Europe, and also enable Member States to substitute an administrative authority for a judicial one in order to fight illegal file sharing.

At the beginning of September the European Data Protection Supervisor (EDPS), which is an independent european authority, published its opinion, which confirmed the analysis of the NGOs.The EDPS was particularly concerned with some amendments which pave the way for the graduated response (or 3 strikes and you're out). This mechanism extends to disputes related to file sharing measures intended fight against terrorism or child pornography.

The EDPS recommended such dispositions to be deleted. The rapporteurs did not follow this primary advices of the EDPS, preferring to try to re-write those amendments criticised, so as to limit their effect. There was some progress, but it must be said that the re-writing of the amendments gave rise to a rather vague, loose text, which introduced concepts that were unknown to European law, and which were taken directly from the proposals of the French cinema lobby (like the "cooperation" between ISPs and producers).

The primary risk is that this rather vague text might be used by certain Member States to give permission to administrative authorities to restrict, without any prior judicial decision, the freedom of expression and information of internet users accused of unauthorized copying.

This risk is real. In July the French government proposed a bill transferring repressive power to an administrative authority which would act at the request of producers of content. It may be voted in November. The United Kingdom also wishes to take the same steps.

In our opinion, the European Parliament must eliminate this risk, which could question both the principle of proportionality and the separation of power, but which could also weaken the acceptability of those measures which are necessary to fight crime.

It must not be possible to extend the measures that national public authorities can implement to fight terrorism or child pornography to disputes concerning non-profit sharing of music and film on the Internet between individuals. Internet users exchanging works without permission should not be treated in the same way as criminals.

This is why we ask you to vote for Amendment 138 to the Trautmann report, tabled by a wide spectrum of MEPs (Guy Bono, Daniel Cohn-Bendit, Zuzana Roithova, Michel Rocard, Marielle de Sarnez, Christofer Fjellner, Rebecca Harms, Marco Cappato, Jean-Luc Benahmias and others).

Amendment 138 states that the national regulation authorities will ensure that no restriction concerning freedom of expression and information of a citizen is taken wthout a prior decision of the judicial authority, except in cases of 'force majeure', threats to security or national criminal law.

Amendment 138 is a guarantee that a bill like the french one about graduated response will not be adopted in Europe. This amendment is in the line with the Bono resolution adopted in April by the European Parliament.

We also ask you to vote against Amendment 34 to the Harbour report which would allow Member States to take measures which harm privacy. This amendment puts national security, crime and file sharing on the same level !

There are other zones which are not at all clear. As well as voting for Amendment 138 and against Amendment 34, we also invite you to clarify the rest of the text. Please find attached all our recommendations concerning the vote on the Telecoms Package.

We hope that you will feel concerned by our request and thank you for your time and attention.

Yours faithfully,

Graduated Response: The Lesson - October 7, 2008

The European Commission opposed on Monday a flat refusal to French president Nicolas Sarkozy's request for deleting amendment 138 of the Telecoms Package. It is yet another slap in the face for the proponents of the graduated response.

Amendment 138, adopted on Sept. 24th by a wide majority of the European Parliament, directly opposes the French draft law setting up graduated response, as Nicolas Sarkozy explained himself to the Commission.

European Commission spokesman Martin Selmayr said: "*The European Commission respects this democratic decision of the European Parliament. In our opinion this amendment is an important re-affirmation of the basic principles of the rule of law in the EU, in particular the fundamental rights of its citizens.*"[13]

La Quadrature du Net welcomes this decision and this confirmation of its own analysis: French draft law is contrary to Community Law, especially because it sets up an administrative authority (HADOPI) entitled to cut internet access off in the name of Copyright. Yet, except when threats to public security are at stake, such a restriction can only be taken by the judicial authority, as reminded by amendment 138.

Thus, French government must go back to the drawing board, especially because, along with being harmful to fundamental rights, its draft law won't add an extra eurocent to artist's revenues.

"*Graduated response*": Will France disconnect Europe? - November 1, 2008

On Wednesday 29 and Thursday 30 October, the French "*Creation and Internet*" law implementing the "*graduated response*" or "*three strikes approach*" against filesharers was passed by the French Senate. Under the watchful eye of the lobbyists who campaigned for it and are the sole beneficiaries of this law (including Vivendi, and representatives of the French cinema and music industries[14]), this expensive project launched personally by President Sarkozy in November 2007, was adopted without any opposition. It was rushed through in a very short time which allowed little opportunity for debate, considering that the liberty of citizens in the digital environment is at stake. The French

[13] http://www.euronews.net/fr/article/06/10/2008/sarkozy-urges-rejection-of-internet-amendment
[14] SACD, SACEM, SNEP, etc.

Senate thus positions itself in total contradiction with European Parliament.

Graduated response: French democracy serving obsolete industrial interests

The law enables the introduction of three-strikes measures[15] against file-sharers and Internet users. The French Senate went along with Nicolas Sarkozy's line, a group of elderly statesmen and women, clinging like crusty old fossils to an economic model from our industrial past - and thus clearly against the European Parliament
The Senators were legislating on an area which they knew nothing about[16]. They had to decode, word by word, the arguments laid out for them by the industries who would benefit, and by the Culture Minister, Christine Albanel.

"*Inconsistencies, lies, confusion and insults which the creative industries habitually use to blame their clients served as justification for a hurried vote, which ignored the wider public debate which is taking place in France and in Europe.*" summarized Jérémie Zimmermann, co-founder of *La Quadrature du Net*.

The vote was carried unanimously in record time, with no problems, surprises or any significant opposition and only around 20 Senators present. However, the law may not get such an easy ride when it goes to the National Assembly, which is the next stage of its emergency process through the French legislature.

The more we learn about the practical implementation of graduated response, the more it becomes obvious that it inherently cannot function without a large-scale surveillance system on the Internet. The technical processes that the Senators were asked to believe in as workable solutions, are in fact, easily circumvented. They will inevitably mean large-scale sanctions against internet users who have done nothing wrong and do not infringe copyright (false positives) and will not bother much those users who do download copyright-protected material (false negatives). The right to defence for people who are accused under the terms of the law does not exist, because neither their innocence nor their guilt can be conclusively proved[17]. The law is therefore already anti-constitutional.

A blind denial of European Parliament's concerns.

The rushed vote of the Senate therefore totally negates the Bono/Cohn-Bendit/Roithova amendment voted in the European Parliament by 88% of its members, on September 24th. Amendment 138 to the Telecoms Package explicitly states that only the judicial authority can impose restrictions on citizens' fundamental rights and freedoms. The French "*graduated response*" totally opposes to this fundamental principle, allowing for an administrative body to arbitrarily deprive citizens from their lives and activities in the digital environment.

"*Such a blatant lack of respect to the European lawmaking process is badly justified by the minister Albanel with a misleading reasoning, questioning whether access to Internet is a fundamental right. This is totally irrelevant. The fundamental principle at stakes is that every citizen has the right to be judged in a fair trial.*" explained Gérald Sédrati-Dinet, analyst for *La Quadrature du Net*.

Moreover, the minister Albanel seemed confident about a removal by the Council of Europe of amendement 138[18]. France has already proven that it was using its presidency to pressure the Commission and the Council in that direction.

"*This law is a scandal, and only serves those industries who refused to move with the times. The Senate is helping the French government to stamp on European democracy, and deny the fundamental rights of citizens. These maneuvers represent the worst side of politics and are a terrible way to write the law. This is how Member States are deepening the democratic deficit in European Union. We must bring it to a stop, for France and for the rest of Europe*" said La Quadrature's representatives.

[15] "*In a nutshell, under such types of schemes ("graduated response" or "3 strikes approach"), copyright holders would identify alleged copyright infringement by engaging in systematic monitoring of Internet users activities. After identifying Internet users alleged to be engaged in copyright violation by collecting their IP addresses, copyright holders would send the IP addresses of those alleged to be engaged in copyright violation to the Internet Service Provider who would warn the subscriber to whom the IP address belongs about his potential engagement in copyright infringement. Being warned by the ISP three times would result in the ISPs termination of the subscribers Internet connection.*", European Data Protection Supervisor (EDPS) comments, 2 September 2008: http://www.laquadrature.net/wiki/EDPS_Comments_Telecoms_Package_IMCO

[16] The sole exception was Bruno Retailleau, representing the Economic Affairs committee, who made a few valiant efforts to intervene, and proposed a fiscal penalty or fine as an alternative to termination of Internet access.

[17] The procedures leading to internet access cut are based on IP addresses listings...

[18] Albanel declared that, since European Commission has qualified this amendment as a recall to some fundamental legal principles, "*this amendment has no legal scope. [...] But meanwhile, it has created an interference effect, a potential manipulation effect, and indeed that is why we hope that this amendment will be withdrawn and we have good reasons to believe that it will be, and, for the sake of clarity, this seems important to me.*"

Commission accepts amendment 138 against graduated response - November 7, 2008

The European Commission accepts amendment 138 (Bono/Cohn-Bendit/Roithova) against the french "*graduated response*", one week after the French law is unanimously voted in first reading by the French Senate.

http://europa.eu/rapid/pressReleasesAction.do?reference=IP/08/1661&format=HTML&aged=0&language=EN&guiLanguage=en
in its plenary vote on 24 September provides that "no restriction may be imposed on the fundamental rights and freedoms of end-users, without a prior ruling by the judicial authorities, notably in accordance with Article 11 of the Charter of Fundamental Rights of the European Union on freedom of expression and information, save when public security is threatened where the ruling may be subsequent."

As already stated on 6 October, the Commission can accept this amendment, which was voted by a nine-tenths majority in the European Parliament plenary (MEMO/08/681).

The Commission considers this amendment to be an important restatement of key legal principles of the Community legal order, especially of citizens' fundamental rights. It leaves Member States sufficient scope for reaching a fair balance between different fundamental rights, in particular the right to respect for private life, the right to protection of property, the right to an effective remedy and the right to freedom of expression and information.

La Quadrature du Net: "*Mr Minister, ...*" - November 20, 2008

Translation of the letter sent by La Quadrature to the French minister in charge of the Telecoms Package, Luc Chatel.

Mr Minister,

On November 27th, the Council of the European Union will examine the project reforming electronic communications, also known as "*Telecoms Package*", as amended by the European Parliament in its first reading last September 24th.

Protection of fundamental rights of European citizens using internet has become one of the major issues at stake of this law proposal.

After lengthy debates in the referred committees, and after the intervention of the EDPS, the independant European authority in charge of the protection of personal data, the European Parliament adopted a series of amendments to the Commission proposal. The MEPs wanted to guarantee that the current level of protection of the european citizens will be at least maintained by the Member States in the future.

But the main safeguard introduced by the Parliament amendment 138 adopted by 88% of the MEPs might be removed by the Council on November 27th, following a request from the French government. The French economic newspaper La Tribune has announced that French government has already managed to convince every other Member State to refrain from voting or to vote in favor of the removal of this amendment.

Nevertheless, as the European Commission underlined in an official memo, this amendment is "*an important restatement of key legal principles of the Community legal order, especially of citizens' fundamental rights. It leaves Member States sufficient scope for reaching a fair balance between different fundamental rights, in particular the right to respect for private life, the right to protection of property, the right to an effective remedy and the right to freedom of expression and information.*"

In its memo, the Commision stated that it would not ask for its removal, contrary to Nicolas Sarkozy's request to the President of the Commission.

The only reason for France to request the removal of this amendment is that it is opposed head on to the French law proposal "*Creation and Internet*" that aims at creating a special court for the Internet users whose account has been used to make unauthorised copies of music and movies1. It is also for France about legalising a posteriori an administrative decision authorising private companies to carry out some police missions on Internet, thus opposing the European policy on personal data.

Therefore, we ask you to oppose the removal of amendment 138 in order to respect, as the European Commision has done, the democratic vote of the European Parliament who has insisted on underlining that fundamental democratic principles, such as the principle of separation of powers or the principle

of proportionality, also apply on Internet, at a time where the Member State assuming the presidency of the European Union seems to have forgotten it.

Failing which, everyone might assess your commitment in the construction of an Europe that protects the fundamental rights of its citizens and the reality of European Democracy.

Hoping that you will be able to act upon this issue, citizenly yours,

Citizen safeguards striked out in EU Council - November 26, 2008

The EU Council reached a political agreement on the telecommunication reform ("*Telecoms Package*") on Thursday, Nov. 27th. On one hand, crucial modifications to the text finally doom Nicolas Sarkozy's project to impose graduated response to the whole Europe. On the other hand, important safeguards to citizen's fundamental rights and freedoms were deleted. The agreed text lowers the protection of privacy in the EU, in the name of "*security*".

During last weeks, citizens from many European countries[19] raised awareness of their ministers representatives in Council on the Telecoms Package, by meeting them, sending letters, alerting the press, etc. This intense activity undoubtedly helped modifying critical parts of the text agreed by the ministers of the twenty-seven Member States.

The dispositions imposing the "*graduated response*" scheme in the European Union (or "*three strikes and you're out*") were neutralized in the Council's version. This is a striking blow at the entertainment industries who spent much effort in promoting it.

But the agreed text contains major problems:

- An unacceptable revision has been made to Art.6, par.6 of ePrivacy directive[20], allowing private operators to collect and process traffic data and exonerate from current privacy rules. This exoneration is so broad that it allows any web company to process any citizen's[21] data for broad, undefined purposes and for a potentially infinite duration. This is a frightening decline, openly ignoring the European Data Protection Supervisor (EDPS)'s recommendations[22] that were followed by the European Parliament.
- The Bono/Cohn-Bendit/Roithova amendment 138 (Art.8, par.4 (ga) of the Framework directive), restating an essential principle of protection of fundamental rights and freedom in European law, has been removed. Deletion of amendment 138, on the vague pretext that the wording was too broad, is in opposition with the democratic expression of 88% of the Members of European Parliament (MEPs) and the European Commission. It clearly shows from the Council a disturbing lack of political courage in protecting citizens' fundamental rights and Freedom
- The dispositions related to "*cooperation in the promotion of lawful content*" (Art.33 par.2a of Universal Service directive), still included in the text, must be deleted as well as the related recitals. These texts were pushed by the French cinema lobby to introduce the graduated response. Their vague wording could still be used by national governments to violate fundamental rights of their citizens until a European court remind them the law. This is exactly what France aims to do in few weeks with its "*graduated response*" law.

"*It is not tolerable that the Council deleted crucial implementations of fundamental principles. The Council lets Nicolas Sarkozy free to violate French citizen's rights with his national project of graduated response. The Commissioner, Viviane Reding, has been quite alone not to share this Pilate's position. The Commission even recalled today*[23] *that this project is a freedom-killer.*" explained Gérald Sédrati-Dinet, analyst for *La Quadrature du Net*.

[19] Letters were sent and published from Czech Republic, France, Germany, Netherlands, Poland, Portugal, Spain, United Kingdom, etc.

[20] Formerly amendment 181 adopted by European Parliament.

[21] Whether they are clients from this operator or not

[22] "*Because [the provision] is broadly constructed, for example, [...] it does not limit the type of entities (data controllers) it is meant to apply, the EDPS is concerned that it could be interpreted too broadly. In particular, the EDPS is concerned that it could be used to legitimise the collection of traffic data for purposes that are not purely security related. He is also concerned that it could open the door for anyone, not only providers of security services and products, to process traffic data alleging to do it for security purposes.*": http://www.laquadrature.net/wiki/EDPS_Comments_Telecoms_Package_IMCO

[23] French newspaper La Tribune published today the Commission response to France's obligation of notification for its "*graduated response*" law

"*The Council's agreement failed to protect fundamental rights by deleting two important safeguards. Let's hope the European Parliament will fight against industry lobbies during the second reading and finally clean up the whole package. No compromise must be made in preserving the right to a due process and privacy in the digital environnement.*" concluded Jérémie Zimmermann, co-founder and coordinator of La Quadrature.

La Quadrature du Net publishes its answers to the Green Paper on Copyright in the Knowledge Economy - November 26, 2008

La Quadrature du Net has submitted on 26 November 2008 its comments[24] on the Green Paper on Copyright in the Knowledge Economy published by the European Commission. Our comments and answers[25] to questions contain important constructive recommendations for the future of the European copyright and author rights framework.

Comments from *La Quadrature du Net* on the Green Paper on Copyright in the Knowledge Economy

General comments on the Green Paper and our recommendations on General issues

We praise the European Commission for having opened a wide-ranging consultation on copyright in the knowledge economy. We particularly welcome the following aspects aspects of the Green paper:

- The recognition that the various facets of knowledge constitute one interdependent ensemble, where science, education, culture, public expression and innovation contribute together to a knowledge society.
- The openness demonstrated by the European Commission with regard to the possible creation of new exceptions, thus showing readiness to reconsider the exhaustive character of the list of possible exceptions in directive 2001/29/CE.
- The recognition that the creation of a new making available to the public exclusive right and more generally the stronger definition and enforcement for exclusive rights have not benefited authors at large.
- The mention that the list of questions are of an indicative nature and that comments can be formulated on other issues relevant to the scope of the Green Paper.

However, some elements in the Green paper indicate that the awareness of the drawbacks of the approach implemented in the past 10 years is still too limited. Sentences such as: A high level of copyright protection is crucial for intellectual creation or A rigorous and effective system for the protection of copyright and related rights is necessary to provide authors and producers with a reward for their creative efforts and to encourage producers and publishers to invest in creative works are of a purely declarative nature. What we mean by this is that these sentences risk to hide the most important question: Which system of copyright protection is likely to serve the aims of rewarding creators at large, of ensuring investment in a wide variety of creative works, and of enabling an empowering access to knowledge and culture? This is all the more surprising since the Green paper acknowledges that many categories of authors and performers do not think that the present system is effective from these view points. Furthermore, for science and research that constitute an important part of the scope of the Green paper, investment in producing creative works is not done by publishers of copyrighted works, who only invest in their dissemination and promotion.

We also point the European Commission to the fact that in the Internet era, the delineation of activities can not be defined by the nature of institutions conducting them. Education happens also outside of the limits of teaching organizations such as schools and universities. Even research related-activities need to made possible also outside of research organizations, in particular for the sake of creating a more productive interface between science and society. This has a bearing on our recommendations for research and education (see below).

General approach to exceptions and limitations

The approach to exceptions implemented in directive 2001/29/CE is a clear policy failure. The exhaustive character of the list of possible exceptions was at the time of its adoption intended to give legal certainty to IPR holders in order to facilitate the adoption of an increased set of exceptions favourable to access and usage of knowledge. This adoption has not happened in practice, or only to a

[24] http://www.laquadrature.net/files/LQdNcommentsonCopyrightGreenPaper.pdf
[25] http://www.laquadrature.net/files/LQdNcommentsonCopyrightGreenPaper.pdf

very limited degree in some Member States. The exhaustive character of the list of exceptions is now standing as an absurd constraint, unjustified by the overall international legal framework. It risks to hinder the putting in place of alternative remuneration schemes based on collective licensing for the non-market exchange of creative works over the internet, at least in situations where legal licensing would be necessary to overcome the opposition of some entrenched and inefficient oligopolies. These schemes are today one of the key paths towards the creation of a sphere of free cultural exchanges and the development of a rich creative economy. We provide below a number of recommendations that aim at re-opening the policy space so that the challenges of creating a knowledge society can be addressed. These recommendations are not limitative, and we also point the European Commission to the approach to Exceptions and Limitations proposed in the Draft Treaty on Access to Knowledge[26]

Recommendation 1: Table a proposal to remove the exhaustive character of the list of exceptions in 2001/29/CE and make clear that new exceptions and limitations can be created as long as they respect the applicable international legal framework (three-step test when applicable, also taking in account other facilities that are open by the Bern Convention Appendix or article 40 of TRIPS, for instance).

Recommendation 2: Propose Member States for the European Union to adopt an open approach to the creation of an instrument of Limitations and Exceptions, going beyond the present work in WIPO on exceptions for the disabled by addressing also minimal research and education exceptions, for instance.

Recommendation 3: More generally, promote a reasonable interpretation of the three-step test (along the line of the declaration A Balanced Interpretation of the Three-Step Test in Copyright Law[27]) in the relevant international arenas (WIPO, WTO) and adopt it for the evolution of the European copyright framework.

Recommendation 4: Oppose the inclusion in trade agreements being negotiated such as ACTA (or other international agreements) of any provision that could directly or indirectly further limit the existing or possible exceptions, or otherwise restrict directly or indirectly the rights of users of knowledge in its widest sense.

***

The Study on the Implementation and Effect in Member States' Laws of Directive 2001/29/EC on the Harmonisation of Certain Aspects of Copyright and Related Rights in the Information Society[28] has shown that the approach based on non-mandatory exceptions has failed in two respects: it has not led to harmonization, but in effect to counter-harmonization, and the possibility of implementing new exceptions to compensate for putting in place TPMs protected against circumvention has not been significantly used in most Member States.

Recommandation 5 in answer to questions 3, 4, 5: We encourage the European Commission to propose that for an additional set of exceptions to be made mandatory. This should be the case in particular for research and education, and for a right of quotation applicable to all media and whose extent is defined in relation to the purpose of the new expressive or creative work that makes use of quotation. We urge it to clarify that mandatory exceptions must be effective in the face of technical protection measures, that is a TPM that does not enable the full exercise of a mandatory exception can be legally circumvented for the purpose of exercising the exception.

Encouragement for an efficient overall contractual management of rights, enabling wide and effective user rights

Whereas (18) of 2001/29/CE, as well as the creation of new extended collective licenses in various countries, have confirmed their validity in European law. However, these schemes, as well as other forms of mechanisms for globally managing rights in a manner that does not create transaction costs or harm to freedoms have not been sufficiently considered in copyright-related policy proposals.

Recommendation 6 in answer to question 2: We call the European Commission to stress the potential of extended collective licenses for non-commercial peer-to-peer exchange between individuals of digital works on the Internet as a possible strategy for ensuring an effective remuneration and funding of creation in a manner that is compatible with the rights and freedoms of all. We call the European Commission to encourage experimentation of such schemes that are growingly considered by collective management societies in Europe.

[26] Draft treaty on Access to Knowledge, 9 May 2005, http://www.cptech.org/a2k/a2k_treaty_may9.pdf

[27] http://www.ip.mpg.de/shared/data/pdf/declaration_three_steps.pdf

[28] L. Guibault, G. Westkamp, T. Rieber-Mohn, P.B. Hugenholtz, (et al.), Study on the Implementation and Effect in Member States' Laws of Directive 2001/29/EC on the Harmonisation of Certain Aspects of Copyright and Related Rights in the Information Society, report to the European Commission, DG Internal Market, February 2007, http://www.ivir.nl/publications/guibault/Infosoc_report_2007.pdf

Specific issues

Licensing with publishers cannot compensate for the absence of a needed exception

Licensing of copyrighted works for activities such as commercial publishing of course will always occur in research, educational or cultural organizations. However, examples in Member States, for instance France, have amply demonstrated that an approach based solely on licensing from publishers totally fails to create the needed environment for access to and use of knowledge resources in research, education, libraries, archives, and museums. On the contrary, it distorts competition by creating preferential access to some resources rather than to others. It enables publishers to gain control and sometimes access to confidential information on usage in educational or research organisations. Finally, it leads to unacceptable restrictions on the type of usage that is authorized, where it can take place, and which public can benefit from it. We advise the European Commission to abstain from further exploring this path (voluntary licensing from publishers). Where a true exception for research and education is in place, for instance in the Scandinavian countries, the results are visible in the output of education and research. When an exception is in place, such as for access by the disabled, voluntary agreements can be useful to install favourable technical conditions for an effective access by the disabled, including in the frame of creative and knowledge creation activities. However, schemes for such agreements must apply to all publishers, which is probably easier to attain through regulatory provisions (see next section).

Recommendation 7 in answer to questions 7, 13 and 19: Only when adequate exceptions for the access to and use of knowledge are in place can licensing agreements between publishers and knowledge organizations (caring for the needs of the general public or of specific groups such as the disabled) play a mutually beneficial role.

Mandatory provisions on access formats for the disabled should specify properties of formats and not specific formats

Recommendation 8 in answer to question 14: mandatory provisions on formats in which works are made accessible for the disabled are useful and even necessary in order to guarantee an effective access. However, these provisions, like any other provision regarding technology or formats, should not mandate usage of specific formats, but rather define properties of the usable formats. The needed properties are being open standards in the sense of the European Interoperability Framework developed by IDABC[29], availability of simply free / open source software solutions for accessing and processing the format, the fact that the format is adapted for access by the disabled, as well as for re-use.

Rather than trying to marginally repair the harm from the directive 96/9/EC on the legal protection of databases, the European Commission should propose for it to be repelled

The 96/9/EC constitutes the prototype of legislative failure. The study in its impact conducted for the European Commission has concluded that it had adverse impacts on access to knowledge and no documentable positive impact on the knowledge publishing industry as a whole. The directive principles are rejected even in countries that are generally favourable to the extension of IPR scope and enforcement. Despite these findings, the Commission has not seriously considered the only convincing option in such a situation which is to repel the directive. This has become a stand case for democracy: is the European Union able to correct one of its mistakes (any government or political institution is likely to make some)?

Recommendation 9 in answer to question 18: We urge the European Commission to face this problem without eye-blinds and to propose Member States to repel the directive. It would do a lot for the standing of the European Union as a supporter of knowledge societies.

Mandatory minimum rules for education and research

Recommendation 10 in answer to questions 22 and 23: We support a mandatory exception for research and education, where the definition of beneficiaries is focused on activities rather than on nature of institutions. However, it could be useful to clarify for instance that educational organizations are by nature beneficiaries of the exception for all their educational activities, provided this applies regardless of the targeted public (for instance it should also apply to open universities or courses open to the general public). We support for the mandatory exception to apply for both education and research as these activities are often inseparable, and even more as the case for a mandatory exception is equally convincing for both.

In contrast, we do not think that mandatory rules on the length of excerpts of works which can be reproduced or made available for teaching and research purposes are to be the preferred scheme.

[29] http://ec.europa.eu/idabc/servlets/Doc?id=19529
A similar definition is included in article 4 of the french Loi n2004-575 du 21 juin 2004 pour la confiance dans l'économie numérique, http://www.legifrance.gouv.fr/affichTexte.do?cidTexte=JORFTEXT000000801164

Whether a given use of works enters in the exception depends on the needs of the teaching activity (as we already pointed when discussing the right of quotation). The only case where a minimum length of extracts would be useful is when defining compulsory rules for any technical protection measure (meaning that a TPM would be illegal in case it does implement the ability to freely extract up to that length). However, this can turn to be harmful if such a rule is interpreted as defining a standard for normally authorized activities. It is for judiciary authorities and accumulating case law to judge if a given use respects the defined scope of the exception. The role of legislation is to provide a clear definition of this scope and to make sure that the exception can be exerted in practice.

User-created content

Let us first remark that all creative or knowledge works are user-created content. However, we understand that the European Commission has here in mind the generalisation of content production by individuals re-using some existing content, and also of the ability of these end-users to reach for the general public. This generalization is one of the most promising developments of the information and knowledge society.

Recommendation 11 in answer to questions 24 and 25: We support the introduction of rules defining or restating acts that users are authorized to accomplish when make use of copyrighted material in their productions, as well as their duties in this respect, provided:

- that these rules never limit the general rights of users such as the right of quotation for the sake of criticism, review or public political expression, and more generally any re-use right that contributes to freedom of expression
- that requirements on duties such as attribution do not introduce harmful technical or human complexity in their implementation. We encourage the European Commission to follow the good practice of Creative Commons licenses and of free re-use licenses in this respect.

We recommend that actions to foster legal certainty for user-generated content activities are conducted with a primary focus on enabling users to conducting these activities. The decision on making or not a specific exception must first consider what can be achieved by way of general exceptions and other user rights.

Internet & Creation by Philippe Aigrain (co-founder of La Quadrature)

Internet & Creation[30] affirms the right for all individuals to share cultural works that have been published in digital form between themselves without profit, in a non-market sphere of exchanges. Where this sharing is stigmatized as piracy by some, the author describes it as a long recognized right that has now become possible to implement at a much greater scale. This is an object of enthusiasm, but calls for an an adapted framework of implementation. Meanwhile, at the end of a period where mass cultural industry have briefly dominated the distribution of cultural works, the material interests of artists and producers of knowledge at large are very poorly served. The rights that were defined for their benefit have now been captured by a few large corporations who maximize their profit on each work by limiting the number of works that will in practice be exposed to the attention of the public. As such limitation is almost impossible to preserve in the internet age, they intend to turn the Internet into something else, a new channel for centrally controlled distribution of consumer works. They will of course fail, but much harm can be done in the process of this failure.

Internet & Creation defines a complete framework for putting in place a mechanism to give us all the best of the internet potential for culture. This framework consists of:

- A precise definition for the non-market sharing of digitally published works that it proposes to recognize. The precise definition makes sure that the channels that provide the greatest part of remuneration to creation will not be harmed by peer-to-peer exchange.
- Putting in place a creative contribution that will be paid by all internet broadband subscribers. A framework is proposed on how to define the amount of this fee. Its product would be used half for the remuneration of works that have been shared over the Internet, and half for the funding of the production of works and the creation of an environment for their dissemination and quality recognition by all. The level of remuneration aims at guaranteeing that the creators will not be negatively impact by the recognition of sharing.

Internet & Creation discusses all aspects of the overall proposal:

- Its legal basis as a licensing to end-users (in contrast to other proposals of licensing catalogs to distributors or ISPs)
- The setting and evolution of the creative contribution
- How the economy of the production of works can remain balanced, including for works such as movies
- The governance of the distribution of the remuneration and of the funding to creation
- International aspects in situations where the proposal would be implemented initially only in some parts of the world
- A non-intrusive (in privacy) usage observation for the remuneration, based on the a large panel of voluntary Internet users and statistical techniques to make it resistant to fraud and efficient to measure the usage of works of lower (but still deserving) popularity
- Paths towards putting in place the proposal.

A key element in the book proposal is to install a situation where contributors to creation and users of works work together for a common good: culture and its sharing by all. Every reader can now act to turn it into a reality.

[30] Internet & Creation has been published in French by Philippe Aigrain at Editions In Libro Veritas in October 2008, http://www.ilv-edition.com/librairie/internet_et_creation.html. An English translation is in the works to be published in the 1st semester of 2009.

Sample press review

Herald Tribune

PARIS: Prodded by the music industry and government, some Internet service providers are reluctantly exploring the adoption of an old-fashioned shunning ritual as the ultimate 21st century punishment: banishing errant online users. (...)

Lawmakers in the European Parliament, in a symbolic vote Thursday, expressed their opposition to the three-strikes approach, which has been championed by President Nicolas Sarkozy of France and explored by governments of other countries, from Britain to Japan to Australia. Many consumer groups are also ghting such proposals, and at least one British service provider is promising a rebellion.

It s a breach of our civil liberties, said Christofer Fjellner, a Swedish legislator in the European Parliament who sponsored the measure, an amendment to a report on cultural industries in Europe. *When government limits access to the Internet it s like limiting freedom of speech. It s like banning people from printing books.* (...)

We believe it s a threat, particularly to public liberties, said Christophe Espem, co-founder of a French group, Squaring the Net, formed to challenge proposed Internet restrictions. He noted that the new administrative authority that would rule on o enses would be outside the legal system.

http://www.iht.com/articles/2008/04/13/business/ISP14.php?page=1

Euractiv

Behind this unusual rejection lies the LIBE Committee s intention to allow the processing of electronic tra c data by *any natural or legal person* , without the consent of the user, if it is necessary for security purposes. Socialist and Green MEPs belonging to the IMCO Committee are not at ease with this wording.

Tra c data include several pieces of information which are considered private by many, particularly IP addresses (the rst source of identity in the online world) and information relating to the duration, timing, volume and origins of an electronic communication. (...)

Civil liberty group Squaring the Net says the LIBE Committee amendment represents *a major breach for the protection the protection of personal data and privacy, as it allows businesses to remotely control users electronic communications without their consent* . It adds that such a measure *paves the way for the deployment of intrusive technologies on the client* .

http://www.euractiv.com/en/infosociety/eu-parliament-split-electronic-data-protection/article-174108

Le Monde

Les operateurs de telecommunications et les fournisseurs de services Internet s abriteraient, ecrit la SACD, *derriere les larges exonerations de responsabilite* des directives relatives au commerce electronique et la vie privee *pour s abstenir de toute action de lutte contre les atteintes aux droits d auteur sur les reseaux numeriques* . La SACD entend ainsi pro ter du *paquet telecom* pour *revenir sur cette defaillance du cadre communautaire* , en y introduisant *des dispositions visant ameliorer le respect des droits d auteur* (...)

La Quadrature du Net, un collectif qui s etait fait connatre en publiant le projet de loi dit *Olivennes* sur la lutte contre le telechargement illicite, et qui a depuis publie la note de la SACD, quali e de *cavaliers legislatifs* les amendements en question, qui reviendraient *abaisser, au nom de la protection de la propriete intellectuelle , le niveau de protection des donnees personnelles et de la vie privee* .

ArstTechnica

The EU Parliament voted Wednesday to pass the "*Telecom Package,*" (...)

Hundreds of amendments were tabled, making the entire legislative process difficult to follow, but two of the key changes proposed were Amendments 133 and 138. As the UK's Open Rights Group points out, 133 would have prevented EU countries from requiring local ISPs to filter content.

138, introduced by a French Socialist MEP Guy Bono (who gets extra points in our book for that moustache) would have prevented any action against Internet users without prior judicial intervention. In other words, Bono insisted that courts need to be involved in any disconnection procedure–exactly the sort of slow process backers of graduated response plans hope to avoid.

http://arstechnica.com/news.ars/post/20080928-eu-parliament-judges-must-be-involved-in-three-strikes-rules .html

La Stampa

La Commissione europea ha risposto con un netto "*no*" oggi a Bruxelles alla richiesta del presidente Sarkozy di bocciare la posizione dell'Europarlamento contro un progetto di legge francese per la protezione della propriet intellettuale su internet. La nuova legge, che dovrebbe essere approvata entro l'anno, attribuirebbe all'authority di regolazione nazionale delle Tlc il potere di monitorare il traffico via internet e tagliare il collegamento alla rete di presunti 'pirat scoperti a scaricare abusivamente film, musica e qualunque altro contenuto coperto dal diritto d'autore. Il 24 settembre scorso, nell'ambito del pi generale 'pacchetto telecom', il Parlamento europeo ha approvato una sorta di 'censura preventiv nei riguardi di questo progetto di legge, approvando a grandissima maggioranza (573 voti contro 74) l'emendamento 138 in cui si chiede ai regolatori nazionali di applicare il principio secondo il quale, salvo che in caso di minaccia per la pubblica sicurezza, nessuna restrizione pu essere imposta sui diritti e le libert fondamentali degli utenti finali, senza la previa autorizzazione delle autorit giudiziarie, segnatamente in accordo con l'Art.

http://www.lastampa.it/_web/cmstp/tmplrubriche/tecnologia/grubrica.asp?ID_blog=30&ID_articolo=5200

Heute

Franzsischer Senat beschliet umstrittenes Antipiraterie-Gesetz.

Das Gesetz sieht ein abgestuftes Sanktionssystem gegen die Nutzer von Tauschbrsen vor. Wer beim Herunterladen urheberrechtlich geschtzter Werke erwischt wird, wird zunchst per E-Mail verwarnt. Wird er innerhalb von sechs Monaten ein zweites Mal ertappt, kommt die nchste Verwarnung als Einschreiben per Post. Tauschbrsennutzern, die sich Musik und Filme trotz der beiden Verwarnungen weiterhin illegal aus dem Netz besorgen, soll schlielich der Internetzugang gekappt werden. Dauer der Netzsperre: von einem Monat bis zu einem ganzen Jahr (...)

Whrend die franzsische Unterhaltungsindustrie den neuen Gesetzentwurf feiert, kommt von Brgerrechtlern und Verbraucherschtzern heftige Kritik. Das Gesetz sei ein Skandal und diene einzig den Interessen der Unterhaltungsindustrie, erklrte die franzsische Brgerrechtsorganisation "*La Quadrature du Net*" ("*Die Quadratur des Netzes*"). Es fhre zu mehr berwachung, beschneide wichtige Grundrechte und missachte zudem eindeutige Entscheidungen des Europischen Parlaments. Tatschlich hatte sich das Europische Parlament noch Mitte September deutlich gegen das franzsische Drei-Stufen-Modell zur Bekmpfung der Internetpiraterie ausgesprochen.

http://www.heute.de/ZDFheute/inhalt/30/0,3672,7399838,00.html

more: http://www.laquadrature.net/press-review

Political Memory

Front page

http://www.laquadrature.net/wiki/Political_Memory

Reach members of European Parliament and track their votes & opinions!

Political Memory is a toolbox designed to help you reach members of European Parliament (MEPs), and track their voting records. We hope it will help citizens to get to better know their elected representatives, and to allow them to inform them on the issues covered by La Quadrature du Net.

We strongly encourage you to improve these informations in order to store MEPs declarations and opinions about these issues.

All theses pages have been produced and maintained using a free software framework, described here.

Your MEPs by country

Austria	Belgium	Bulgaria	Cyprus
Czech Republic	Germany	Denmark	Estonia
Spain	Finland	France	United Kingdom
Greece	Hungary	Ireland	Italy
Lithuania	Luxembourg	Latvia	Malta
Netherlands	Poland	Portugal	Romania
Sweden	Slovenia	Slovakia	

List of recorded votes

Click each directive,report or resolution for more informations and sorting of the MEPs according to their votes.

- Directives reforming the EU's regulatory framework for electronic communications networks and services (Telecoms package), September 24th 2008
- Rapport Bono on cultural industries in Europe (Bono report), April 10th 2008
- Resolutions on European Patent Litigation Agreement (EPLA), October 12nd 2006
- Directive on patentability of "computer-implemented inventions" (software patents), September 24th 2003

Search MEPs

By political group

- Group of the European People's Party (Christian Democrats) and European Democrats
- Socialist Group in the European Parliament
- Group of the Alliance of Liberals and Democrats for Europe
- Union for Europe of the Nations Group
- Group of the Greens/European Free Alliance
- Confederal Group of the European United Left - Nordic Green Left
- Independence/Democracy Group
- NI Non-attached Members.

Others

- Search MEPs alphabetically, by name
- Search MEPs by committee
- Search MEPs by their office location within the Parliament
- Members of the Conference of Presidents

Telecoms Package score by MEP

http://www.laquadrature.net/wiki/Telecoms_package_directives_1st_reading_by_score

Political Memory: Directives reforming the EU's regulatory framework for electronic communications networks and services, Results by Final Score

- Detailed scores

Rank	MEP	Country	Group	Final score	
1	Jean-Marie LE PEN	France	NI NI	94.1	Show details
2	Marine LE PEN	France	NI NI	94.1	Show details
3	Marco CAPPATO	Italy	ALDE	91.2	Show details
4	Jules MAATEN	Netherlands	ALDE	91.2	Show details
5	Lydia SCHENARDI	France	NI NI	91.2	Show details
6	Margrete AUKEN	Denmark	Verts/ALE	90.9	Show details
7	Heide RÜHLE	Germany	Verts/ALE	90.6	Show details
8	Marie-Hélène AUBERT	France	Verts/ALE	90.0	Show details

Telecoms Package score by country

http://www.laquadrature.net/wiki/Telecoms_package_directives_1st_reading_by_country

Political Memory: Directives reforming the EU's regulatory framework for electronic communications networks and services, Results by Country

- Sweden: mean score: 56.5/100
- Czech Republic: mean score: 55.3/100
- Netherlands: mean score: 54.8/100
- Denmark: mean score: 54.7/100
- Austria: mean score: 54.6/100
- Belgium: mean score: 54.1/100
- Cyprus: mean score: 53.7/100
- Luxembourg: mean score: 52.9/100
- Finland: mean score: 52.6/100
- Italy: mean score: 51.6/100
- Germany: mean score: 51.5/100
- Portugal: mean score: 51.2/100
- France: mean score: 50.8/100
- United Kingdom: mean score: 49.9/100
- Greece: mean score: 48.8/100
- Latvia: mean score: 48.0/100
- Malta: mean score: 47.9/100
- Spain: mean score: 47.5/100
- Bulgaria: mean score: 46.7/100
- Romania: mean score: 44.3/100
- Hungary: mean score: 44.1/100
- Poland: mean score: 43.2/100
- Estonia: mean score: 42.4/100
- Slovakia: mean score: 42.0/100
- Slovenia: mean score: 41.0/100
- Lithuania: mean score: 40.9/100
- Ireland: mean score: 36.2/100

MEP page

http://www.laquadrature.net/wiki/JanellyFourtou

Political Memory: Janelly FOURTOU, MEP

General Data

- *Born on* 04 February 1939, Paris
- *Country*: France
- *Political Group*: Group of the Alliance of Liberals and Democrats for Europe (ALDE)

  Party: Avenir Démocrate

Contact

- Parlement européen
  Bâtiment Altiero Spinelli 09G142
  60, rue Wiertz, 1047 Bruxelles
  Tel.: +32 2 28 45 150/+32 2 28 47 150
  Fax: +32 2 28 49 150
- Parlement européen
  Bâtiment Winston Churchill M02072
  1, avenue du Président Robert Schuman - CS 91024, 67070 Strasbourg Cedex
  Tel.: +33 3 88 175 150/+33 3 88 177 150
  Fax: +33 3 88 179 150
- *Email*: janelly.fourtou(at)europarl.europa.eu
- *Page on European Parliament website*: www.europarl.europa.eu/members/expert/alphaOrder/view.do?language=EN&id=4336

Functions in European Parliament

- Committee on Petitions (Member)
- Committee on the Internal Market and Consumer Protection (Member)
- Delegation to the EU-Chile Joint Parliamentary Committee (Member)
- Delegation to the Euro-Latin American Parliamentary Assembly (Member)
- Committee on Legal Affairs (Substitute)
- Delegation for relations with the countries of Central America (Substitute)

Votes

- 24/09/2008 - Directives reforming the EU's regulatory framework for electronic communications networks and services **26.2/100**
- 10/04/2008 - Rapport Bono on cultural industries in Europe **0.0/100**
- 12/10/2006 - Resolutions on European Patent Litigation Agreement (EPLA) **5.6/100**
- 24/09/2003 - Directive on patentability of "computer-implemented inventions" (software patents) **18.9/100**

MEP score

http://www.laquadrature.net/wiki/Telecoms_package_directives_1st_reading_details_by_score?showmep=JanellyFourtou

Political Memory: Directives reforming the EU's regulatory framework for electronic communications networks and services, Results in Detail by Final Score

MEP	Final score	am. 98/2	am. 120	am. 138/1	am. 133/rév.	am. 30/1	am. 34	am. 62/2	am. 62/4	am. 67/2	am. 75/2	am. 76/2	am. 112	am. 117	ams. 155+172	am. 193
Janelly FOURTOU	26.2	for	against	against	against	for	for	for	for	for	for	for	for	for	against	for

Show details for all MEPs

- am. 98/2: Imposition of restrictions on access to services and applications that endangers net neutrality - recommendation: against, coefficient: 1
- am. 120: Deletion of relation with EUCD and IPRED which is out-of-scope - recommendation: for, coefficient: 1
- am. 138/1: Obligation for a ruling by judicial authorities before applying any restriction to end-users' fundamental rights, except for public security, preventing from 3-strikes approach - recommendation: for, coefficient: 5
- am. 133/rév.: No technology mandate for Internet filtering - recommendation: for, coefficient: 4

Neusprech im Überwachungsstaat

Martin Haase

Die Anschläge vom 11. September 2001 machten den Bereich der „Inneren Sicherheit“ wieder zu einem wichtigen Thema der Politik. Die meisten gesetzlichen Neuregelungen schränken die bürgerlichen Freiheiten ein und sind daher sehr umstritten: man denke an die Gesetzesänderungen, die unter der Bezeichnung „Großer Lauschangriff“ bekannt wurden, an die Einführung biometrischer Ausweise, die kontaktlos ausgelesen werden können (der so genannte E-Pass oder der elektronische Personalausweis), die Vorratsdatenspeicherung, die heimliche Online-Durchsuchung und -Überwachung oder die Ausweitung der Videoüberwachung. Gerade in jüngerer Zeit erregte das neue, stark umstrittene BKA-Gesetz für Aufregung: Die Aufhebung einer klaren Grenze zwischen geheimdienstlichen und polizeilichen Ermittlungen, die Möglichkeit, ohne richterliche Genehmigung geheim zu ermitteln, die fehlende parlamentarische Kontrolle der geheimdienstähnlichen Aktivitäten des BKA und die Schwächung der Bundesanwaltschaft können als Bedrohung grundsätzlicher Rechte angesehen werden, weshalb damit zu rechnen ist, dass das Gesetz in der derzeitigen Form vor dem Bundesverfassungsgerichts scheitern wird. Auch Maßnahmen, die gar nicht in den Bereich der Sicherheitspolitik gehören, erweisen sich als problematisch zumindest aus datenschutzrechtlicher Perspektive: die Gesundheitskarte durch die damit verbundene Zentralisierung von Versichertendaten, die Steueridentifikationsnummer, die ebenfalls mit einer zentralen Datenerfassung verbunden ist, oder auch die Registrierung von Nummernschildern durch Polizeikameras und Mautbrücken.

Solche Maßnahmen sind der Bevölkerung schwer zu vermitteln, daher müssen Politiker, die sich für sie einsetzen, zu außergewöhnlichen sprachlichen und rhetorischen Mitteln greifen, um die negativen Auswirkungen solcher Gesetze zu verschleiern oder klein zu reden und ihnen positive Seiten abzugewinnen.

1 Neusprech 1.0

In George Orwells berühmten dystopischen Roman *1984* [8], der vor sechzig Jahren geschrieben wurde, stellt der Autor einen Überwachungsstaat vor, der sich einer besonderen Sprache bedient, die abweichendes Verhalten undenkbar macht, weil mehrdeutige Wörter ausgemerzt wurden und das Vokabular so organisiert ist, dass staatliche Maßnahmen möglichst positiv und alternativlos erscheinen. Die neue Sprache heißt *Newspeak* oder in deutscher Übersetzung *Neusprech*. Dass es Orwell um zeitgenössische Sprachkri-

tik geht (besonders mit Bezug auf die Politik), zeigt sein Aufsatz *Politics and the English Language* [10].

Das Vokabular der neuen Sprache umfasst drei Bereiche, wie sie Orwell in einem Anhang zu seinem Buch beschreibt:

A Dieser Bereich umfasst einen vereinfachten Alltagswortschatz. Unregelmäßige Formen sind ausgemerzt: *think* bedeutet ‚denken' und ‚Gedanke' (das unregelmäßige Nomen *thought* ist verschwunden). Ein Wortbildungssystem, das an Esperanto erinnert, ermöglicht es, aus wenigen Grundwörtern weitere abzuleiten: so kann mit dem Suffix *-ful* ein Adjektiv aus einem Nomen abgeleitet werden: *thinkful* bedeutet somit ‚gedanklich', ‚gutes Denken' (im Sinne der Staatsdoktrin) ist *goodthink*, das Gegenteil *ungoodthink*, das sich mit *plus* und *doubleplus* verstärken lässt; eine Person oder Sache, die durch ganz böse Gedanken auffällt, kann also als *doubleplusungoodthinkful* qualifiziert werden. Orwell orientiert sich hier an der Plansprache *Basic English*, der er kritisch gegenübersteht; allerdings fließt auch seine Kenntnis des Esperanto-Wortbildungssystems hier ein, denn während eines Parisaufenthaltes lebte Orwell in einem esperantosprachigen Haushalt [4].

B Dieser Bereich des Wortschatzes umfasst politisches Vokabular. Zur Verschleierung der Inhalte politischer Doktrinen oder Maßnahmen werden die Wörter aus gekürzten Wortstämmen zusammengesetzt und orthografisch angepasst. Beispiele sind: *ingsoc* (die Staatsdoktrin, die etymologisch auf *English socialism* zurückzuführen ist), *artsem* (für *artificial insemination*, die staatlich verordnete Form der Fortpflanzung), *crimethink* (‚Gedankenverbrechen') und *minitrue* (für *Ministry of Truth*, das ‚Wahrheitsministerium' organisiert die Staatspropaganda und schützt die Bevölkerung vor *crimethink*).

C Zusätzlich zu den genannten Bereichen des Wortschatzes ist es erlaubt, wissenschaftliche Terminologie zu verwenden; es sind natürlich nur Wissenschaften erlaubt, die nicht im Widerspruch zur Staatsdoktrin stehen, weshalb *ingsoc* auch die Bedeutung ‚Wissenschaft' hat.

Im Zusammenhang mit der Sprache der inneren Sicherheit ist der B-Wortschatz von besonderem Interesse. Es ist nicht auszuschließen, dass sich Orwell hier von Viktor Klemperer hat inspirieren lassen: Sein *Notizbuch eines Philologen* über die Sprache des Nationalsozialismus' erschien 1947 [6]. Klemperer beschreibt, wie im Nationalsozialismus neue Wörter geprägt wurden, die sich zum Teil bis heute erhalten haben wie *Staatsakt* oder *Dachorganisation*; zudem benennt er typische Verfahren der nationalsozialistischen Wortbildung, in der die Wortbildungselemente *Groß-* und *Volks-* reihenbildend werden. Der Sprachkritiker entlarvt Floskeln wie: *in stolzer Trauer*, die in Todesanzeigen (insbesondere von gefallenen Soldaten) häufig war, nach Ende des Nationalsozialismus' aber wieder verschwand. Zudem zeigt er, wie im Nationalsozialismus oft verwendete Wörter bedeutungsentleert werden und sich in ihrer Wirkung abschwächen; so zum Beispiel: *fanatisch, historisch* und *ewig*.

Als Begründer moderner PR-Techniken kann Edward Bernays angesehen werden. Sein 1928 erschienenes Buch *Propaganda* [1], zielt nicht nur auf politische Propaganda ab,

sondern betrifft vor allem den Bereich der *public relations* in Wirtschaftsunternehmen. Natürlich wurden Bernays' Vorschläge auch von den Nationalsozialisten befolgt, finden sich in Orwells Newspeak wieder und haben bis heute nichts an Aktualität eingebußt. Dass sich die Sprache der Politik und der Massenmedien an Bernays' Vorschlägen orientiert, zeigte Noam Chomsky bereits in einem Vortrag von 1984 [2], aus dem später ein Buch und sogar ein Film entstanden, beide unter dem Titel *Manufacturing Consent* [3].

2 Neusprech 2.0

Auch die Neuauflage politischer PR-Sprache in Bernays's Sinn, die wir heute erleben, kann als Neusprech bezeichnet werden. Wieder geht es darum, negativ-besetzte Wörter durch eine positive oder wenigstens undeutliche Entsprechung zu ersetzen. Der nächste Abschnitt enthält ein paar Beispiele.

2.1 Neue Wörter und Bedeutungen

Die neue allgemeine Krankenversicherungskarte, die der Bevölkerung und ihren Entscheidern trotz aller Datenschutzproblematik schmackhaft gemacht werden soll, heißt im besten PR-Jargon *Gesundheitskarte*. Das Verfahren, einen Sachverhalt (hier: *Krankheit*) durch sein Gegenteil zu beschreiben (hier: *Gesundheit*) wird in der Rhetorik als *Antiphrase* bezeichnet. Es geht in diesem besonderen Fall vor allem darum, etwas Negatives und möglicherweise Tabuisiertes durch etwas Positives zu bezeichnen. Man kann daher auch von *Euphemismus* sprechen.

Gerade der Bereich der Sicherheitspolitik enthält eine Reihe von interessanten Euphemismen: Das Prinzip, Verdächtige zu erschießen, das man in anderen Zusammenhängen als *Schießbefehl* bezeichnet hätte, wird *gezieltes Töten* genannt oder pseudo-englisch *Targeted Killing* [12]; gelegentlich wird auch von *finalem Rettungsschuss* gesprochen. Innenminister Schäuble fordert: „Gefährder zu behandeln wie Kombattanten und zu internieren“ [13]; *internieren* ist ein Euphemismus für ‚wegsperren‘, ‚behandeln‘ ist aufgrund der Bedeutungsvielfalt des Wortes (*Polysemie*) ein Euphemismus, da unklar bleibt, was mit den so genannten *Gefährdern* geschieht. Während *Kombattant* ein juristisch definierter Terminus ist, der sich auf Teilnehmer an Kriegshandlungen bezieht, ist völlig unklar, was mit *Gefährder* gemeint ist. Es handelt sich um einen juristisch (bisher) undefinierten Begriff, der so etwas wie einen (möglicherweise) Verdächtigen oder zukünftigen Störer bezeichnet; in einem Rechtsstaat ist jemand, der noch kein Verbrechen begangen hat, jedoch immer unschuldig und kann nicht einfach weggesperrt werden. Die Gleichsetzung eines *möglichen* Störers der inneren Sicherheit mit einem Teilnehmer an Kriegshandlungen verwischt zudem den Aufgabenbereich der Polizei mit dem des Militärs.

Es handelt sich ebenfalls um einen Euphemismus, wenn Wahlcomputer als *Wahlmaschinen* bezeichnet werden, denn bei Maschinen denkt man weniger an Manipulation als bei Computern. Der niederländische Wahlcomputerhersteller Sdu ging sogar soweit, sein computerbasiertes Wahlsystem *NewVote* [14] zu nennen, um positive Assoziationen zu wecken, wohl ohne daran zu denken, dass diese Bezeichnung sehr an Orwells *Newspeak* erinnert.

Auch der Terminus *Vorratsdatenspeicherung* ist ein Euphemismus, denn *Vorrat* ist ein positiv besetztes Wort und *Datenspeicherung* ist eine relativ neutrale Umschreibung; der englische Terminus ist hingegen negativ konnotiert: *data retention* erinnert an Gefängnis oder zumindest an die Einschränkung von Freiheit. Treffender wäre es, von *Datenhortung* zu sprechen. Auch *Online-Durchsuchung* ist eine Bezeichnung, die offenbar positiv wirken soll: ähnlich wie das *E-* in *E-Pass* vermittelt *online* ein Gefühl von Modernität; *Durchsuchung* täuscht vor, dass es sich um eine Maßnahme handelt, die mit einer Hausdurchsuchung vergleichbar ist. Es handelt aber vielmehr um eine *Überwachungsmaßnahme*, die vielleicht treffender als *Computerverwanzung* bezeichnet wird – dann natürlich mit einem Wort, das negative Assoziationen erweckt. Üblich ist auch die Bezeichnung *Bundestrojaner*, da das auf den Computer eingeschleuste Spähprogramm als *Trojanisches Pferd* oder *Trojaner* bezeichnet wird.

Mithilfe von Metaphern gelingt es, bestimmte vermeintlich positive Aspekte hervorzuheben und negative zu verschleiern: Wenn zum Beispiel gefordert wird, der „genetische Fingerabdruck“ solle als biometrisches Merkmal eingesetzt werden, so werden Merkmale des persönlichen Erbinformation mit dem in der Kriminalistik erprobten Fingerabdruck gleichgesetzt, obwohl die eigene DNA sehr viel persönlichere Informationen enthält als ein Fingerabdruck und die Eindeutigkeit der untersuchten DNA-Merkmale noch nicht zweifelsfrei erwiesen ist. Dass natürlich auch der klassische Fingerabdruck als Mittel einer eindeutigen Identifikation problematisch ist, muss hier nicht vertieft werden.

2.2 Grammatische Aspekte

Allerdings betrifft Neusprech nicht nur den Wortschatz, sondern auch die Grammatik der Politikersprache zeigt Auffälligkeiten: Besonders interessant ist das Pronomen der ersten Person Plural *wir*. Dass dieses Pronomen politisch verwendet werden kann, zeigt schon Jewgenij Samjatins Roman Мы (*wir*) von 1920 [18], den George Orwell 1946 [9] rezensiert und wohl auch als Inspiration für *1984* genutzt hat. In Samjatins Gesellschaft spielt ein allgemein propagiertes Wir-Gefühl eine besondere Rolle.

In vielen Sprachen können verschiedene Bedeutungen bzw. Verwendungsweisen des Pronomens der ersten Person Plural beobachtet werden:

(1) exklusives *wir*: Das Pronomen bezeichnet lediglich den Sprecher und mit ihm verbundene weitere Personen, schließt jedoch den angesprochenen nicht mit ein. Im Französischen kann der Ausschluss des Angesprochenen durch Hinzufügung von *autres* unterstrichen werden: *nous autres* ‚wir (anderen)‘ (so auch im Italienischen: *noi altri*, im Spanischen ist *otros* jedoch fester Bestandteil des Pronomens geworden und kann auch nicht-exklusiv verwendet werden). Im Tok Pisin, der Staatssprache Papua Neuguineas (eine ehemalige Pidginsprache), gibt es ein spezielles exklusives Pronomen, nämlich *mipela*.

(2) inklusives *wir*: Das Pronomen schließt den Angesprochenen mit ein (im Französischen: *nous* oder umgangssprachlich *on*); das Tok Pisin hat hierfür ein spezielles Pronomen, nämlich *yumi*.

(3) extensives *wir*: Bei dieser Verwendungsweise von *wir* sind Personen gemeint, die in mehr oder weniger mittelbarer Beziehung zum Sprecher stehen, wobei der Sprecher aber gar nicht dazu gehört: so sagen Fußballfans: *Gegen Stuttgart haben wir schon zu oft verloren.* Verloren hat allerdings nicht der Sprecher, sondern eine Mannschaft, mit der er sich identifiziert. Ein extremes Beispiel ist die bekannte Zeitungsüberschrift nach der letzten Papstwahl: *Wir sind Papst.*

Hinzu kommt noch das *wir* des Majestäts- und Bescheidenheitsplurals (*pluralis maiestatis/modestiae*), das aber für die weiteren Erörterungen nicht relevant ist.

Die extensive Verwendung von *wir* ist geradezu typisch für die Sprache von Politikern. Das folgende Beispiel aus einer Rede von Angela Merkel auf einer Wahlkampfveranstaltung der CDU in Osnabrück [16] zeigt, wie das *wir* in unterschiedlicher Funktion eingesetzt wird:

> „Die CDU wird sich nicht davon abbringen lassen. Sie werden sich erinnern – die Älteren unter Ihnen –, wie viele Schlachten *wir* [CDU] schon geschlagen haben: die Videoüberwachung – gestritten mit den Sozialdemokraten. Heute hätten *wir* [extensiv] weder die libanesischen Kofferbomber gefunden, noch hätten *wir* [extensiv] die Schlägereien des alten Mannes in der U-Bahn in München so schnell aufklären können, und heute findet jeder Videoüberwachung auf großen Plätzen, öffentlichen Plätzen, ganz normal. Wenn es die Union nicht gewesen wäre, die dafür gekämpft hätte, dass das notwendig ist, hätten *wir* [Deutschen] heute noch keine Videoüberwachung, und deshalb werden *wir* [CDU] auch andere Themen auf die Tagesordnung bringen, wie bestimmte Veränderungen im Jugendstrafrecht, genauso wie die Online-Durchsuchung, und vieles andere mehr. *Wir* [CDU/Regierung] werden nicht zulassen, dass technisch manches möglich ist, aber der Staat es nicht nutzt – dafür aber die Verbrecher und Täter und Terroristen es nutzen. Das ist nicht *unser Staat* [Deutschland]. Der Staat muss [wehrhaft sein].“

Wir können davon ausgehen, dass Angela Merkel weder die Kofferbomber persönlich gestellt hat, noch an der Aufklärung der „Schlägereien des alten Mannes“ direkt beteiligt war. Das extensive *wir* soll die eigentlichen Ermittler lediglich mit der Politik der Kanzlerin in Verbindung bringen.

2.3 Rhetorik

Über Wortschatz und Grammatik hinaus gibt es eine Reihe von rhetorischen Besonderheiten der Politikersprache: stereotype Metaphern und Floskeln spielen eine besondere Rolle. Seit einigen Jahren spielen Fußballmetaphern eine besondere Rolle. So ist die Partei Silvio Berlusconis nach dem italienischen Fußball-Schlachtruf *Forza Italia!* (‚Kraft [*d. h.:* los!] Italien‘) benannt. Oft ist von einem politischen *Eigentor* die Rede, die Legislaturperiode wird in *Halbzeiten* eingeteilt, und auf Franz Müntefering glaubt: „Der Parteivorsitzende ist der Spielführer.“ [7].

In einem Interview mit Deutschlandradio über das „Internet als Tatmittel der Zukunft“ [17] bemüht Jörg Ziercke, der Präsident des Bundeskriminalamts, zahlreiche Floskeln,

die schon an seiner Intonation (*Allegrosprechweise*) als solche erkennbar sind: so ist Terrorismus immer „international“, Kriminelle sind „skrupellos“, die Kriminalität „schwer“, „schwerst“, „sozialschädlich“ oder wenigstens „organisiert“; „Netzwerke“ (Neologismus für *Netze*) werden „zerschlagen“ und Ziercke will „Beweise verdichten“ (transitiv!) und „im Internet ... Durchsuchung durchführen“, wobei es doch schon reichen würde, im Internet zu suchen. Der BKA-Präsident will „mit dem technischen Fortschritt Schritt halten“ und wird nicht müde, darauf hinzuweisen, dass „99,9 % der Menschen in Deutschland“ von der Online-Überwachung gar nicht betroffen seien, was jedoch bedeutet, dass 80 000 Menschen, also eine mittelgroße Stadt, betroffen sein werden. Norbert Geis, der rechtspolitische Sprecher der CDU/CSU-Fraktion, hatte zuvor von „99 % gesprochen, was 800 000 Menschen ins Visier der Fahnder rücken ließe. Hier kann von einer missglückten Floskel gesprochen werden.

Als Neusprech kann wohl auch angesehen werden, dass Ziercke mit Bezug auf den Terrorismus von „Szene“ und das neue Wort „*Anfasser* für Ermittlungen“ prägt. Hierbei handelt es sich offensichtlich um ein neues Wort, das einen Anfangsverdacht bezeichnet, der eine Ermittlung erst erlaubt.

Ein beliebtes rhetorisches Mittel, schwierigen Fragen auszuweichen, ist das Ablenkungsmanöver. Statt eine Frage zu beantworten, wird eine Anekdote erzählt. Auf die Frage: „Wie können Sie garantieren, dass die Online-Durchsuchung tatsächlich nur in Einzelfällen eingesetzt wird?äntwortet der Innenminister auf einer Pressekonferenz wie folgt [15]:

> „... genehmigen von den G-10-Kommissionen, dann muss im Einzelfall begründet werden. Es gibt mehr Maßnahmen, das ist schon wahr [unverständlich], das will ich nicht sagen. Da haben wir doch, neulich hat die Polizei, Bundesanwaltschaft drei so, die terroristische Anschläge von erheblicher Qualität geplant haben, 600 Kilogramm Sprengstoff ist ja nun keine Kleinigkeit, und die benützen, die haben ungefähr ein Dutzend Handys benützt, nicht, so mit Prepaid-Karten und ein Mal telefonieren kurz, gleich wieder wegschmeißen, nächste. Und dann muss natürlich die Polizei, wenn sie überwachen will, oder der Richter, genehmigt der zwölf Handys, plötzlich tat sich's verzwölffachen. Und deswegen bleibt, der Norbert Geis, dass 99, die Aussage, dass 99 % nicht betroffen sind, wenn er gesagt hätte 99,9 % werden niemals davon betroffen sein, hätte er auch Recht gehabt. Das ist so.“

Es wird eine Anekdote erzählt, die mit der gestellten Fragen in keinem Zusammenhang Zusammenhang steht, denn hier geht es gar nicht um die Online-Durchsuchung, sondern um Telefonüberwachung; zudem zeigt die Anekdote gerade nicht, dass oder wie sichergestellt werden kann, dass die Telefonüberwachung (wie wohl dann auch die Online-Durchsuchung?) nur in Einzelfällen angewandt wird.

Nach der Anekdote wird der Bundesinnenminister jedoch konkreter und kommt auf den (im eigentlich Gesetz aufgegebenen) Richtervorbehalt, um dann mit der rhetorischen Frage zu kontern: „Glauben Sie, das werden viele Maßnahmen sein?“ Danach argumentiert er *ad personam* bzw. *ad ignorantiam* („Leute [...], die keine Ahnung haben“, „ich

versteh nix davon“) und führt schließlich als Autorität den Chef des Bundeskriminalamts an:

> „Das ist gesetzlich sicherzustellen, nur, wir haben ja den, wir haben ja den Vorschlag gemacht, es darf nur in engen Fällen, Abwehr wirklich terroristischer Bedrohung ernsthafter Art, dann muss, kann das Bundeskriminalamt, wir haben sogar gesagt, nur der Präsident darf den Antrag stellen und ein Richter muss ihn genehmigen und er muss begründen warum. So. Glauben Sie, das werden viele Maßnahmen sein? Erstens. Zweitens, um [unverständlich] was immer man unter Online-Durchsuchung versteht, da reden ja auch die Leute alle ganz klug, die keine Ahnung haben. Es ist so aufwändig, dass der Chef des Bundeskriminalamts, der Herr Ziercke, der versteht e bissel was davon. Ich versteh nix davon. Er hat gesagt, so 'ne Maßnahme ist so aufwändig, wir wären überhaupt nur in der Lage, zehn pro Jahr überhaupt zu versuchen, ob sie gelingen, ist noch was anderes. Also der Norbert Geis hatte mit anderen Worten einfach Recht.“

Auch in Schäubles Redebeitrag werden wieder die ominösen 99(,9) % ins Feld geführt, die auf Ausführungen von Norbert Geis zurückgehen; zum Schluss ist jedoch von nur zehn Fällen pro Jahr die Rede. Man kann sich des Eindrucks nicht erwehren, dass hier vor lauter Über- und Untertreibungen (*Hyperbel* und *Litotes* in der Rhetorik) im Trüben gefischt wird, da die Zahlen nicht zusammenpassen. Die Zahlen sind hier auch nicht als Zahlen zu verstehen, sondern als rhetorische Mittel; gleichzeitig vermindert der Redner seine Regresspflicht durch die Verwendung des Konjunktivs II („wir wären... “).

Mit dem vielzitierten rechtspolitischen Sprecher der Unions-Bundestagsfraktion Norbert Geis (CSU) führte Dirk-Oliver Heckmann im Deutschlandfunk am 9. Juli 2007 ein Interview [5], in dem Geis folgende Ausführungen macht:

> „... Ich fordere eine Diskussion darüber, einmal, wie wir umgehen mit der gezielten Tötung eines potenziellen Aggressors und wie wir das gesetzlich, grundgesetzlich absichern können. Darüber fordere ich eine Diskussion. Wenn Sie aber ganz konkret fragen, ich bin natürlich dafür, dass wir einen potenziellen Aggressor, einen Terroristen, der unser Land bedroht, dass wir den natürlich liquidieren können müssen, sonst setzen wir uns unnötig unter Umständen einem Anschlag aus. Wenn ich dadurch Anschläge verhindern kann, muss es möglich sein, solche Anschläge schon im präventiven Bereich abzuwehren.“

Hier spricht Geis von einem „potenziellen Aggressor“, der zu „liquidieren“ (*Euphemismus*) sei; im weiteren Verlauf wird das Wort *Aggressor*, das eher in den Wortschatz des Krieges gehört, durch das bekannte unklare *Gefährder* ersetzt, was mit einer Umschreibung (*Periphrase*) vorbereitet wird: „einen potenziellen Aggressor, einen Terroristen, der unser Land bedroht“. Da das Adjektiv *potenziell* (wie übrigens auch *präventiv*) zum Ausdruck bringt, dass die zu liquidierende Person noch gar nicht zum Täter geworden ist, geht es hier rein rechtlich um eine Tötung von Unschuldigen (man beachte auch die

Doppelung: „im präventiven Bereich abzuwehren“). Geis verwendet nicht nur ein extensives *wir*, sondern ersetzt es am Ende durch ein noch extensiveres *ich*: „Wenn ich dadurch Anschläge verhindern kann. . . “.

Der Interviewer fühlt sich von Geis an das US-amerikanische Gefangenenlager Guantanamo erinnert, und Geis fährt fort:

> „Wir brauchen kein Guantanamo, sondern wir wollen ja nur die potenziellen Gefährder in so etwas wie Unterbindungsgewahrsam, den wir ja jetzt auch schon haben, natürlich unter engen rechtlichen Voraussetzungen. Wir können Hooligans abfangen und können zunächst einmal die Freiheit entziehen, wir können sie also unterbringen, so lange, bis beispielsweise ein Fußballspiel vorbei ist.“

Hier geht es nicht nur um *Gefährder* (also potenzielle Täter), sondern um einen „potenziellen Gefährder“, womit wahrscheinlich ein sehr großer Kreis von Menschen gemeint ist – ja eigentlich die Bevölkerungsmehrheit, da ein potenzieller Gefährder ja noch kein Gefährder, also potenzieller Täter ist, und die Bevölkerungsmehrheit bildet die Gruppe der (noch) nicht-potenziellen Täter. Hier wird nun das Mittel des Wegsperrens vorgeschlagen und verglichen („so etwas wie“) mit einem mehr oder weniger bekannten (nicht unumstrittenen) Rechtsmittel, dem Unterbindungsgewahrsam. Hierbei handelt es sich aber um etwas ganz anderes: in Unterbindungsgewahrsam können Störer genommen werden, um die Störung für kurze Zeit (in der Regel bis zu 24 Stunden) zu unterbinden. Geis benutzt zunächst den Euphemismus „die Freiheit entziehen“ und drückt sich schließlich noch euphemistischer aus, indem er von „unterbringen“ spricht (rhetorisch als *Litotes* bezeichnet).

Den Vergleich von Terrorismus und Fußball greift der Interviewer Dirk-Oliver Heckmann gleich auf: „Aber das Fußballspiel, um da mal einzuhaken, Herr Geis, das ist im Antiterrorkampf eigentlich nie zu Ende.“ Darauf antwortet Geis:

> „Das ist richtig. Und deswegen ist dies ja auch rechtlich ein viel schwierigeres Problem, denn wir werden ja auf Dauer jemandem dann die Freiheit entziehen, indem wir ihn in ein Unterbindungsgewahrsam bringen.“

Hier wird wieder euphemistisch von „Freiheit entziehen“ gesprochen, was aber „auf Dauer“ geschehen soll; man kann vielleicht darüber streiten, ob „auf Dauer jemandem dann die Freiheit entziehen“ eine euphemistische Periphrase für die lebenslängliche Inhaftierung von Unschuldigen ist, jedenfalls ist so etwas nicht mit Unterbindungsgewahrsam gemeint, selbst wenn Geis das glauben machen will. Man beachte hier auch die Verwendung des Futurs, das nahe legt, dass dieses Vorgehen für die Zukunft vorgesehen ist.

Der Interviewer wirft dann auch gleich eine Zwischenfrage ein: „Ohne Prozess?“ Als Antwort auf diesen Einwurf entgegnet Geis:

> „In diesem Fall ohne Prozess. Das geht schwer runter, das sage ich Ihnen, das fällt einem nicht leicht vor allen Dingen dann, wenn man, so wie ich auch, immer wieder sagt, dass wir die freiheitlichste Grundordnung haben, die wir

> je hatten. Aber es geht ja uns darum, diese freiheitliche Grundordnung zu schützen. Und deswegen wehren wir uns gegen die Gefährder. Und wir können nicht warten, bis die Gefährder zuschlagen."

Wieder ist von Gefährdern die Rede, wobei Geis seine Vorschläge (Tötung und Wegsperren) mit dem positiv besetzten Verb *sich wehren* zusammengefasst wird. Während das ersten beiden Vorkommen des Pronomens *wir* alle Deutschen miteinschließt, sind am Ende mit *wir* offenbar nur noch Geis und seine Anhänger gemeint, die allerdings zahlreich sind, wenn man bedenkt, dass er von mehr als der Hälfte der Wähler seines Wahlkreises direkt in den Bundestag gewählt wurde. Seine Ausführungen zeigen, dass er mit der Absicht, die freiheitliche Grundordnung zu schützen, diese abschaffen will, denn wenn wichtige Prinzipien wie die Unschuldsvermutung oder das Recht auf Leben abgeschafft werden, kann auch nicht mehr von einer freiheitlichen Grundordnung gesprochen werden.

* * *

Die hier diskutierten Zitate und weiteres Material (zusammengestellt z. B. im Schwarzbuch Politikerzitate [11]) zeigen, dass wir es lexikalisch (in Bezug auf Neuerungen im Wortschatz), grammatisch (vor allem im Bezug auf das Pronomen *wir*) und rhetorisch in der Politik mit Phänomenen zu tun haben, die als Neusprech bezeichnet werden können. In allen Fällen geht es darum, den Bürgern unpopuläre Einschränkungen der persönlichen Freiheiten als unausweichlich, ja nützlich und vorteilhaft zu verkaufen und für die eigene, eigentlich unangemessene Politik zu werben. Die Wörter *verkaufen* und *werben* sind hier bewusst verwendet, denn es handelt sich hier tatsächlich um *public relations*-Maßnahmen. Dass sich gewisse sprachliche Mittel (*Gefährder, 99(,9) %*) bei unterschiedlichen Politikern (parteiübergreifend) wiederfinden, lässt auf eine konzertierte Aktion schließen, die vielleicht nicht im Einzelnen abgesprochen wurde, sondern auf Konsens basiert. Es ist zudem aufschlussreich, dass Kriegswortschatz (Kombattant, Aggressor) im Bereich der inneren Sicherheit eingesetzt wird, was nahelegt, dass eine Abgrenzung der beiden Bereiche verwischt oder gar aufgegeben werden soll. Die ständige Verwendung eines rechtlich unklaren Wortes wie *Gefährder* ist symptomatisch für den Paradigmenwechsel, der eintritt, wenn dieses Wort zum juristischen Begriff werden sollte: es ist der Übergang vom Rechtsstaat zum Überwachungsstaat. Diesen gilt es zu verhindern.

Literatur

[1] Edward Bernays. *Propaganda.* New York: Horace Liveright, 1928.

[2] Noam Chomsky. The Manufacture of Consent (1984). Peck, James (ed.): *The Chomsky Reader.* New York: Pantheon, pp. 123–136, 1987.

[3] Noam Chomsky and Edward S. Herman. *Manufacturing Consent: The Political Economy of the Mass Media.* New York: Pantheon, 1988.

[4] Vikipedio [Esperanto-Wikipedia]. Novparolo. http://eo.wikipedia.org/w/index.php?title=Novparolo&oldid=2011330, 2008. [zuletzt konsultiert am 1. Dezember 2008].

[5] Norbert Geis. Deutschlandfunk-Interview: Geis hält Freiheitsentzug ohne Prozess für gerechtfertigt. http://www.dradio.de/dlf/sendungen/interview_dlf/644540/, 9.7.2007. [zuletzt konsultiert am 1. Dezember 2008].

[6] Viktor Klemperer. *LTI – Lingua Tertii Imperii.Notizbuch eines Philologen.* Leipzig: Reclam, 1947.

[7] Franz Müntefering. SZ-Interview: Auf die Vermittlung der Arbeitslosen konzentrieren. http://www.bundesregierung.de/Content/DE/Interview/2007/11/2007-11-02-m_C3_BCntefering-sz.html, 2.11.2007. [zuletzt konsultiert am 1. Dezember 2008].

[8] George Orwell. *1984.* London: Secker & Warburg, 1949.

[9] George Orwell. Review of Zamyatin: *We*. *Tribune*, 4. Januar 1946. http://www.orwelltoday.com/weorwellreview.shtml [zuletzt konsultiert am 1. Dezember 2008].

[10] George Orwell. Politics and the English Language. *Horizon*, April 1946. http://www.calvinvanhoek.com/articles/2007/04/politics-english-language/ [zuletzt konsultiert am 1. Dezember 2008].

[11] Schwarzbu.ch-Weblog. Politische Zitate. http://schwarzbu.ch/blog/politische-zitate/index.html, 2008. [zuletzt konsultiert am 1. Dezember 2008].

[12] Wolfgang Schäuble. Terrorismusbekämpfung und Innere Sicherheit aus deutscher Sicht. Rede bei der 4. Handelsblatt-Konferenz „Sicherheitspolitik und Verteidigungsindustrie“. http://www.bmi.bund.de/cln_028/nn_662956/Internet/Content/Nachrichten/Reden/2007/07/BM__Handelsblatt__Konferenz_Sicherheitspolitik,templateId=renderPrint.html, 3.7.2007. [zuletzt konsultiert am 1. Dezember 2008].

[13] Wolfgang Schäuble. Spiegel-Interview: Schäuble fordert Handy- und Internetverbot für Terrorverdächtige. http://www.spiegel.de/politik/deutschland/0,1518,493094,00.html, 7.7.2007. [zuletzt konsultiert am 1. Dezember 2008].

[14] Wij vertrouwen stemcomputers niet [Wir vertrauen Wahlcomputern nicht]. Sdu. http://www.wijvertrouwenstemcomputersniet.nl/Sdu, 2007. [zuletzt konsultiert am 1. Dezember 2008].

[15] Udo Vetter. Law Blog: Plötzlich zwölf Handys. http://www.lawblog.de/index.php/archives/2007/10/22/plotzlich-zwolf-handys/, 2007. [zuletzt konsultiert am 1. Dezember 2008].

[16] Udo Vetter. Law Blog: Wo bleibt das Recht? [Angela Merkel bei einem Wahlkampfauftritt in Osnabrück]. http://www.lawblog.de/index.php/archives/2008/01/22/wo-bleibt-das-recht/, 2008. Transkription: http://www.chauvi.de/blog/20080122-149/ [beides zuletzt konsultiert am 1. Dezember 2008].

[17] Jörg Ziercke. Deutschlandradio-Interview: Internet ist das Tatmittel der Zukunft. http://www.dradio.de/dkultur/sendungen/interview/590511/, 6.2.2007. [zuletzt konsultiert am 1. Dezember 2008].

[18] Евгений Замятин. *Мы*. http://az.lib.ru/z/zamjatin_e_i/text_0050.shtml, 1920. [zuletzt konsultiert am 1. Dezember 2008].

The Trust Situation

Why Data Protection Doesn't Protect Much

Dr. phil. des. Sandro Gaycken
University of Stuttgart
sandro.gaycken@philo.uni-stuttgart.de

Abstract:
Informational self-determination is an important ground for individual and societal notions of freedom. It relies on what we know about those surveilling us. However, as surveillance has become too huge a phenomenon to be "known" in any substantial sense anymore, generating informational self-determination has be delegated to data protection, a judicial canon which is designed to provide us with secure knowledge, at least in principle. But as data protection itself has grown a complicated topic, it has in fact only eroded our capacity to make any informed judgements. What we have to rely on in the end is hear-say knowledge and our trust regarding the involved agencies. And that's not good for our informational self-determination at all.

Exciting Times

Surveillancewise, we sure live in exciting times. Especially in Germany. We have warmongering ministers of the interior, suspecting terrorists and child molesters around every other corner. We have our police tuned in on their new paradigm of preemptive crime fighting (the „Neues Interventionsparadigma"). We have the sensors-industry in a goldrush with loads of fancy new devices, mimicking every sense we have or don't have in a most precise manner. They're joined by the IT-industry spawning ever new routines to track their customers. And unsurprinsingly we have the natural conjunction of all these ideas. Mutually reinforcing, they produce a massive increase in technologically mediated surveillance with a preemptive character, surveilling everyone independently of any real suspicion. Hundreds of such devices are currently developed, lending themselves to all sorts of political, security or commercial interests.[1] They enable the police or abstract entities as the „state" or the „law" just as well as the „market" or the boss to watch us, record us, profile us and sort us in highly efficient ways, readying our virtual selves for further processing of whatever sort. A time to worry? Many say no. We might have increased our surveillance. But we are still far from any sort of orwellian dictatorship. Surveillance is in good democratic control by data protection. Anyone suggesting differently is just one of those overly activated activists, who circle frenetically around highly hypothetical or exceptional cases. And the same people saying this tend to add stuff like: „Come on" or „Get real". However, I have my doubts that we (the activists) just play something up here. Actually, I have a lot of doubts.

One argument for instance which I believe to be quite substantial is that, despite the fact that we are not having a dictatorship *right now*, we cannot give any sort of guarantee for this status to remain over the next, say, 50 years. We wouldn't have learned a bit from history if we would.[2] And if we cannot do this, then why should we develop sociotechnical instruments and organizational structures which are only of limited use in a democracy[3], but of optimal use for any dictator interested in eliminating even the mildest type of hypothetical opposition in no time and with very little personnell? Preemptive surveillance technologies and the correlated social organizations have this kind of potential. The technologies enable their operators to recognise any sort of danger very early by what people do, how they communicate, how they inform themselves, how they move and interact. Thus, if you are dictator, simply re-define „danger" in whatever sense suits you and the machines and their operators will return you a list of all the „dangerous" people in no time.[4] In this

1 See http://www.securityresearchmap.de/ for an overview on what is done in Germany.
2 The Dutch philosopher Mark de Vries recently put very nicely: „the only thing we can learn from history is that we don't learn anything from history" (private communication).
3 About the criminological inefficiency, which is a very important point for the whole discussion, see for instance Albrecht (2008) or Gaycken (2009, forthcoming).
4 This has already be noted in the eighties by Gary T. Marx (1988) in his famous judgment on undercover police surveillance, but is equally valid for any sort of preemptive surveillance. And it definitely needs to be repeated.

sense, the sociotechnical structures of the new paradigm of preemptive crime fighting can have a strongly repressive, anti-democratic side-effect. This is not to say that they *produce* any dictatorship or that they *are* somewhat totalitarian. That wouldn't make much sense. But they do lend themselves very easily to any sort of totalitarian purpose and will maximize many of its effects. Thus, if society cannot control such a dangerous side-effect (or totalitarianism) at all times in the future and if that side-effect can turn devastingly against it at some point, it should be abolished in the present. Not the current freedom right now, but the preservation of our future freedom heavily depends on that. We need to replace the current paradigm of surveillance with a more moderate and sustainable kind. Ok, this is one of my earlier arguments.[5] I like repeating it as I deem it important for everyone to know. But it's actually not the one I want to sketch out here. The one I want to write about here makes a different point. While the first argument points to the fact, that preemptive surveillance cannot be controlled from getting into terribly wrong hands at some point in the future, this next argument will stress that surveillance already made us loose control in one very important respect.

The Death of Informational Self-Determination and an Unexpected Murderer

Now what is this thing which we have lost? I think it is our informational self-determination. It's dead. And what's even worse: one of its two murderers is the butler. Data protection has actually helped substantially in killing it. Let me tell you about this (philosophical) murder tale.

To start the story, we have to dive a little into the German formulation of the right of informational self-determination and try to find its essence. Don't worry, it's less tough (or boring) than it sounds. Informational self-determination states that we should always be able to fully know about anything that is known about us someplace else, so we do not have to suspect anything vicious lying around somewhere where we don't want that. Because – and this is already the essence of the informational self-determination – if we would have to suspect anything, however faint, we might decide to be cautious and adjust our behaviour in such a way that we better not do or say anything against anyone anymore. Just to be on the safe side. But being limited in such a way, we would not be very „free“ anymore, would we? We would no longer be able to speak our mind, to think for ourselves, to act for ourselves. And, conclusively, we could not be able to form a free and democratic society anymore. Such a society consists – per definition – of people who are free in this very aforementioned sense. Thus protecting the informational self-determination is an deeply important for anyone who wants to be free and it should be an innate and ongoing concern of any democratic state. The problem now is that this version of informational self-determination has been formulated in the mid-eighties. Back then, surveillance was limited and the gathering and storing of data was still a pretty hard thing to do and thus easily controllable. These days, things are different. Surveillance is becoming more and more omnipresent and data about us are constantly gathered, stored and processed without the slightest hassle. A change in degree, some might say, but in fact one which puts us into quite a different situation. With such complexity around us, how can we uphold our informational self-determination anymore? Surely noone will expect us to know every miniscule technological measure involved anymore, including all possible side-effects, every agent doing something with our data, intentionally or unintentionally, every institution with all its organizational routines and possible failures, and so on. If you still have a life to live, that is just utterly impossible. But then how can we fully know about anything that is known about us someplace else? The answer is: data protection. Data protection is a judicial corpus of rights which tries to warrant our informational self-determination by doing two things. First, it monitors all kinds of data collections and tries to restrain the collectors from gathering too many too individual data which might be able to spawn very personal informations.[6] Secondly, it tries to bring forth and maintain transparency by informing us about new surveillance measures and by enabling us to understand our particular rights to enforce disclosure on the data and informations gathered about

5 See for instance Gaycken (2008).

6 The distinction between data and information is the following: data is everything that is just bits and bytes, not read out and unable to say anything about anyone, information is everything that is put together in understandable, meaningful words (or images or whatever), able to say something about someone.

us. And that's how it works. Critical information are prevented, other data are noted and made accessible, so we can at least *in principle* be able to know everything known about us.

To Know or not to Know

So no problems anymore? No need to worry? Informational self-determination secured? I would say no. Because all data protection has put up is a huge and complicated apparatus of hypothetical „you-could-have-your-informational-self-determination-if-you-only-would-...". But in fact, noone does all those things data protection suggests. Noone informs herself precisely about all the kinds of data collected, all the agents and institutions involved, all the technologies and organizational structures. Noone understands those numerous rights of data protection in any sufficient detail, let alone enforces them with time and money against the myriads of surveillers we meet every day. Most people haven't even understood what data protection is in general or what it does.[7] Instead, what most lay people know about the current state of affairs is this: Everyone with money and power can surveill me almost constantly. There is some thing called data protection, but that's all lawyers, bosses and politicians again, complicated judicial stuff noone understands or takes the pain to get into. Period. This particular and certainly widely shared perception is the reality of how people meet the situation. Lawyers and scientists will now say: „It doesn't have to be like that. I *can* know everything". But that's exactly the point: *you* can. The average guy can not. Lay people instead *have to* rely on the *public* perception of surveillance and of data protection, not on the professional. And this public perception is the relevant and the only relevant ontology for the information self-determination. We can see that immediately. If we do not directly and precisely *know* everything about surveillance (in a scientific, professional sense of knowing), then what we *assume and suspect* about surveillance is the decisive element. It's the only way how we can judge the situation. And how are such assumptions and suspicions built? They are built by the public perception. By what we already believe about the involved actors and by some prominent cases coming to our ears through the media. Thus data protection has actually built us an illusion of knowledge where there actually is none. And it has done so – given the aim of generating secure knowledge for everyone and concludingly mostly for lay people – in the exact counterintuitive way: by adding more complex things to know.

Informational Mood-Dependent-Uncertainty

This plunges us into chaos. The bad kind of chaos, to be sure. Because as the average non-expert has to rely on the public perception and as that is mostly not very precise and informative on details, any following judgement has to be made in rather substantial informational uncertainty. This is the situation technical complexity *and* complex data protection laws and procedures have put us in. And it's a rather troublesome situation. It's not just that people will not have all the information needed for an informed consent. Much worse is in fact the psychological mechanism that how people relate to imprecise information and how they use it in actual judgements is quite angled. Tversky and Kahnemann have shown this long time ago and it's a well-known fact by now.[8] Hear-say-knowledge is not just knowledge. You pick what you like and forget or deny what you don't like and you rather believe people you trust than people you don't trust. So what you know (or what you think you can know in principle) multidimensionally depends directly on the prior opinion you have about the involved agencies, the confidence you have in them. Such a subjective and colorful „trust situation" is the real ground of any scarcely informed decision-making. For informational self-determination, this can immediately be shown to be generally bad. Because trusting those involved on either side of surveillance is asking quite something. To feel entirely secure and safe, to think and behave freely in this exceedingly surveilled world, people have to trust the state, the police, those in power, the secret services, the tax agencies, the banks, the big companies, the rich-and-richer or – in less

7 Funny story one researcher from the center for data protection in Kiel told me: He went on a tour in some company and was asked by the guide what he does for a living. He answers: „I do data protection". The guide answers, surprised: „Oh! That still exists?"

8 See the landmark book by Tversky/Kahnemann (1982).

institutional terms – lawyers, politicians, state authorities, salesmen and bosses. It is immediately clear that this is not exactely any sort of „most-trusted"-list. Quite the contrary. The informational self-determination of average lay people in fact depends on how trustworthy they consider a number of institutions and people which are commonly considered somewhat shady. Surely not by all people, maybe in all generality not even by the majority. But certainly by a huge part of the population and especially by those in difficult situations or with little access to knowledge in the first place, in other words: the socially weaker. And even for those who might still have full faith in all the agencies just mentioned, any sense of „real" informational self-determination must honestly be given up. All we really have is something like an „informational mood-dependent-uncertainty". And that is – to end the tale – why informational self-determination can be considered dead and partially murdered by data protection. Its successor, the informational mood-dependent-uncertainy, on the other hand is a very subjective, little rational and fragile little thingy, easily disturbed and not at all able to carry the heavy burden laid upon its weak and meager shoulders: our freedom.

A Really Real Concern

This is not just a hypothetical thing out of the brain of a philosopher. A number of already existing examples can be cited. People in need of aid such as troubled families or drug addicts already stop seeking such aid as they fear they might be identified and be observed closely henceforth, with ensuing disadvantages in other situations. Informants of the press remain silent as they cannot rely on their anonymity anymore, knowing that this might just not be guaranteed anymore. In Germany, we just had the Telekom scandal shaking that particular trust situation deeply and certainly for some time into the future, hindering the free press substantially in fulfilling its mission. The message many people got from this is that those „big guys" (and – by typefying extension – *any* „big guys") do not play by the rules anyway. They do what they want. So how can any informant to the press trust data protection to protect him in the first place? The same applies to many attestors or accused in court cases. They fear telling details of their cases to their lawyers as they fear that their lawyers might be wiretapped too. Thus many social arrangements needed in a just and democratic society or arranged in solidarity start to crumble. But do we have the right to exclude people in such situations from „felt freedom", from our free community? Are we ourselves free from this sort of exclusion? Maybe tomorrow someone *we* have to worry about might be interested in something we would rather see protected? Considering this, there are only two options. Either the technologies themselves have to be abolished again. Or we should be honest about the consequences, abolish – aloud and publicly – the whole idea of informational self-determination and say good-bye to freedom for all those who do not boast of confidence into lawyers, politicians and rich and powerful people in general.

References

Albrecht, H.-J. (2008): Kosten und Nutzen technisierter Überwachung. In: 1984.exe. Gesellschaftliche, politische und juristische Aspekte moderner Überwachungstechnologien. Hrsg. S. Gaycken, C. Kurz. transcript Verlag, Bielefeld.

Gaycken, S. (2008): Was habe ich eigentlich zu verbergen? In: Datenschleuder 92, Berlin 2008.

Gaycken, S. (2009, forthcoming): Rhetorik und Realität der Überwachung. In: Kontrollverluste. Interventionen gegen Überwachung. Hrsg. von Leipziger Kamera. Initiative gegen Überwachung, Münster.

Marx, G.T. (1988): Undercover. Police Surveillance in America. Berkeley.

Kahneman, *D., Slovic, P., & Tversky, A. (1982)*. Judgment under uncertainty: Heuristics and biases. New York.

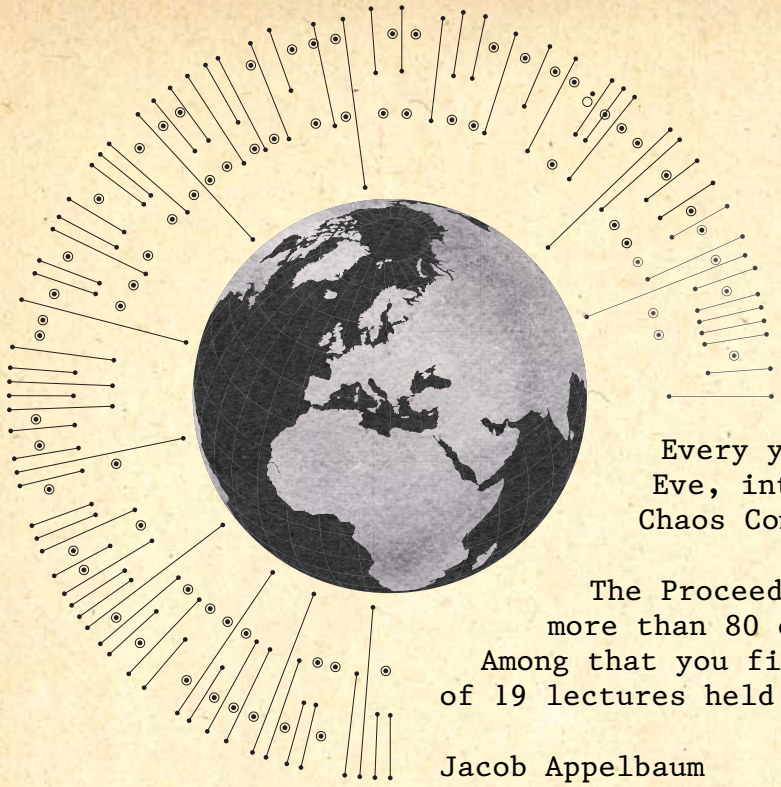

Every year between Christmas and New Years Eve, intergalactic Hacker-society meets up at Chaos Communication Congress.

The Proceedings contain a full description of more than 80 events of this years congress. Among that you find academic papers by the speakers of 19 lectures held on 25C3, here a selection:

Jacob Appelbaum
Advanced memory forensics: The Cold Boot Attacks
Recovering keys and other secrets after power off

wesen
Algorithmic Music in a Box
Doing music with microcontrollers

Tor E. Bjørstad
An introduction to new stream cipher designs
Turning data into line noise and back

Rahmstorf
Climate Change
State of the Science

maha/Martin Haase
Neusprech im Schnüffelstaat
Politikersprache zwischen Orwell und Online

Bernhard Fischer
OnionCat - A TOR-based Anonymous VPN
Building an anonymous Internet within the Internet

Jan Torben
Privacy in the social semantic web
Social networks based on XMPP

Sandro Gaycken
The Trust Situation
Why the idea of data protection
slowly turns out to be defective

ISBN 978-3-934636-06-4 02300
9 783934 636064