Chair of Network Architectures and Services
Faculty of Informatics
Technical University of Munich

Information Session for the Seminars “Future Internet” and “Innovative Internet Technologies and Mobile Communications”

Prof. Dr.-Ing. Georg Carle and I8 research staff
Organization: Daniel Raumer
Contact: seminar@net.in.tum.de

Challenging Topics!
Sometimes previous knowledge required!

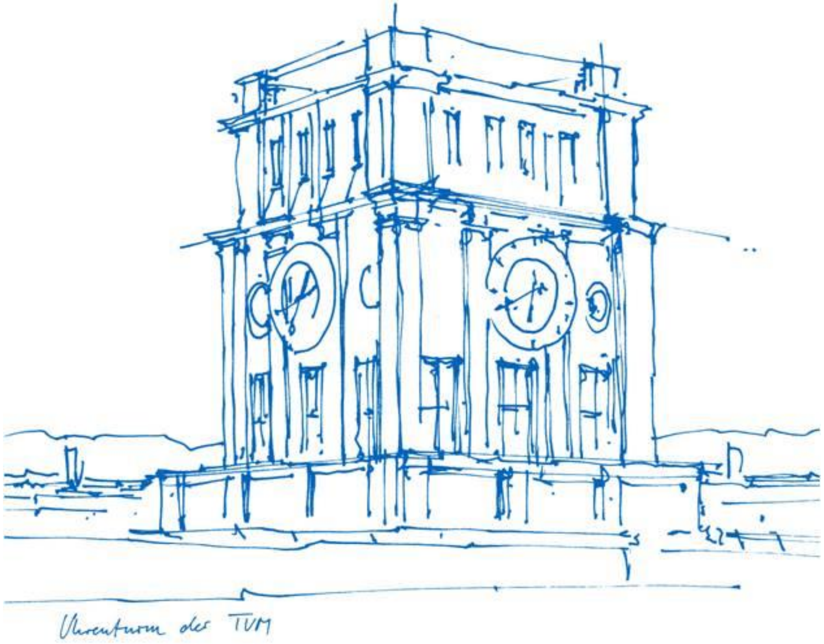

Content

www.fotoila.de

Administrative Things for all Seminars

- Basic Information
- Grading
- Registration
- About the topics

Basic Information

- Lecturer:
  - Prof. Dr.-Ing. Georg Carle

- Organization: seminar@net.in.tum.de (only use this mail address!)
  - Daniel Raumer
  - tba.

- Overview
  - Main Language: German
    - we will offer an English track (presuming a minimum of 4 participants)
  - Extent: 2 SWS (4 ECTS)
    - 4 ECTS * 30 hours = 120 working hours expected from you
  - Course Type:
    - For M.Sc. Students: Master's Seminar (Master-Seminar)
    - For B.Sc. Students: Advanced Seminar Course (Seminar)

- Languages: German and English
  - English only speakers are welcome to join (seminar will be split in two tracks if necessary)

English Only Track

- We offer an English only track if at least one non-German (native) speaker wants to attend the seminar

- The English only track will have separate sessions
  - Probably 1-2 sessions (depending on the number of students)

- Attendance not mandatory for talks in the “standard” track
  - Students in the “standard” track also don’t have to participate in the English track talks
  - You are still welcome to join the other track’s talks ☺

- Usually the English track is quite small
  - This means less attendance (if the opportunity to improve your English is not a good enough incentive for you…)

Paper Procedure

Talks and Papers can be in German or English!

- First version of your paper
  - Agree on the content with your advisor
  - Use the supplied paper template from the webpage
  - Keep in touch with your advisor
  - Try to finish well in time so you advisor can give you feedback
- Write reviews
  - You will be given two papers of your fellow students
- Final version of your paper
  - Use the received reviews to improve your paper
  - You will also receive some feedback from your advisor
  - If you and your advisor agree → publication in the seminar proceedings
- Extent
  - Your paper MUST be 6-8 pages in ACM 2-column style (including references etc.)

Topic Handling

- From your advisor(s) you may receive some literature.
  - This is just to get you started

- Find appropriate (scientific) sources yourself
  - scholar.google.com
  - acm.org
  - ieee.org
  - semanticscholar.org
  - You sources’ sources
  - …

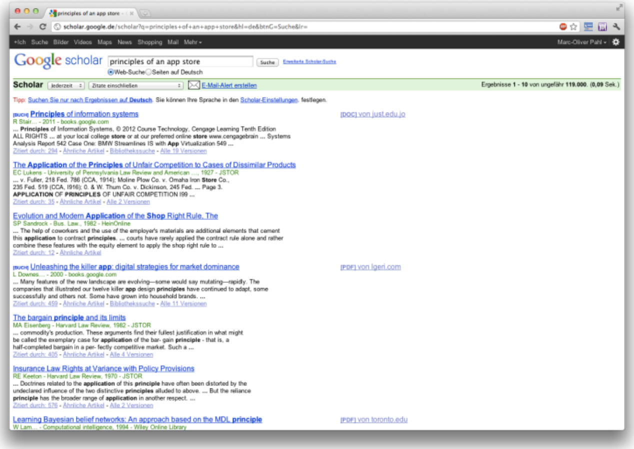

principles of an app store

● Web-Suche ○ Seiten auf Deutsch

Ergebnisse **1 - 10** von ungefähr **119.000**. (**0,09** Sek.)

Tipp: Suchen Sie nur nach Ergebnissen auf **Deutsch**. Sie können Ihre Sprache in den Scholar-Einstellungen. festlegen.

[BUCH] **Principles** of information systems — [DOC] von just.edu.jo
R Stair... - 2011 - books.google.com
... **Principles** of Information Systems, © 2012 Course Technology, Cengage Learning Tenth Edition ALL RIGHTS ... at your local college **store** or at our preferred online **store** www.cengagebrain .... Systems Analysis Report 542 Case One: BMW Streamlines IS with **App** Virtualization 549 ....
Zitiert durch: 294 - Ähnliche Artikel - Bibliothekssuche - Alle 19 Versionen

The **Application** of the **Principles** of Unfair Competition to Cases of Dissimilar Products
EC Lukens - University of Pennsylvania Law Review and American ...., 1927 - JSTOR
... v. Fuller, 218 Fed. 786 (CCA, 1914); Moline Plow Co. v. Omaha Iron **Store** Co., 235 Fed. 519 (CCA, I916); 0. & W. Thum Co. v. Dickinson, 245 Fed. .... Page 3. **APPLICATION** OF **PRINCIPLES** OF UNFAIR COMPETITION I99 ...
Zitiert durch: 35 - Ähnliche Artikel - Alle 2 Versionen

Evolution and Modern **Application** of the **Shop** Right Rule, The
SP Sandrock - Bus. Law., 1982 - HeinOnline
... The help of coworkers and the use of the employer's materials are additional elements that cement this **application** to contract **principles**. ... courts have rarely applied the contract rule alone and rather combine these features with the equity element to apply the shop right rule to ...
Zitiert durch: 12 - Ähnliche Artikel

[BUCH] Unleashing the killer **app**: digital strategies for market dominance — [PDF] von lgeri.com
L Downes... - 2000 - books.google.com
... Many features of the new landscape are evolving—some would say mutating—rapidly. The companies that illustrated our twelve killer **app** design **principles** have continued to adapt, some successfully and others not. Some have grown into household brands. ...
Zitiert durch: 459 - Ähnliche Artikel - Bibliothekssuche - Alle 11 Versionen

The bargain **principle** and its limits
MA Eisenberg - Harvard Law Review, 1982 - JSTOR
... commodity's production. These arguments find their fullest justification in what might be called the exemplary case for **application** of the bar- gain **principle** - that is, a half-completed bargain in a per- fectly competitive market. Such a ...
Zitiert durch: 405 - Ähnliche Artikel - Alle 4 Versionen

Insurance Law Rights at Variance with Policy Provisions
RE Keeton - Harvard Law Review, 1970 - JSTOR
... Doctrines related to the **application** of this **principle** have often been distorted by the undeclared influence of the two distinctive **principles** alluded to above. ... But the reliance **principle** has the broader range of **application** in another respect. ...
Zitiert durch: 576 - Ähnliche Artikel - Alle 2 Versionen

Learning Bayesian belief networks: An approach based on the MDL **principle** — [PDF] von toronto.edu
W Lam... - Computational intelligence, 1994 - Wiley Online Library

Just presenting the given literature is NOT enough

The Advisor’s Role

Advisors created topics within their research context.
→ They have broad knowledge about the context of your seminar topic.

You task is to do research and write a scientific text about a specific topic beyond basic lecture content.
→ Your advisor is not responsible for you tasks.

Adhering to the deadlines is your responsibility.
→ Your advisor will not remind you.

Advisors will help you if you ask them to.
→ Keeping contact with your advisor allows you to write a much better seminar paper.

Advisor can give you feedback
→ Ask for feedback about your first paper version, the peer reviews, your slides for the talk, etc.

Talk Procedure

Talks and Papers can be in German or English!

- Prepare your talk
  - Finished slides must be discussed with advisor 1 week before the talk
  - Advisors usually offer the opportunity of test talks

- Give your talk

- Session chair for one talk
  - Introduce the talk
  - Watch the time constraints
  - Try to get the discussion started after the talk (ask at least one question if nobody else does)

- Mandatory attendance on all sessions in your track
  - If you cannot attend for a good reason contact seminar@net.in.tum.de <u>in advance</u>

Grading

1. Both of your paper submissions (6–8 pages in ACM) (50%)
   - 1st version: 37,5%
   - 2nd version: 12,5%

2. Your talk (20–25min, following discussion and feedback) (25%)
   - Content is graded
   - Personal presentation style is not

3. Your reviews of papers from other seminar participants (25%)

Improving Your Presentation Skills

You have the chance to get your talk recorded

- Have a **look at yourself** after the talk!

- Your talk was great? Share it and show it to your friends.

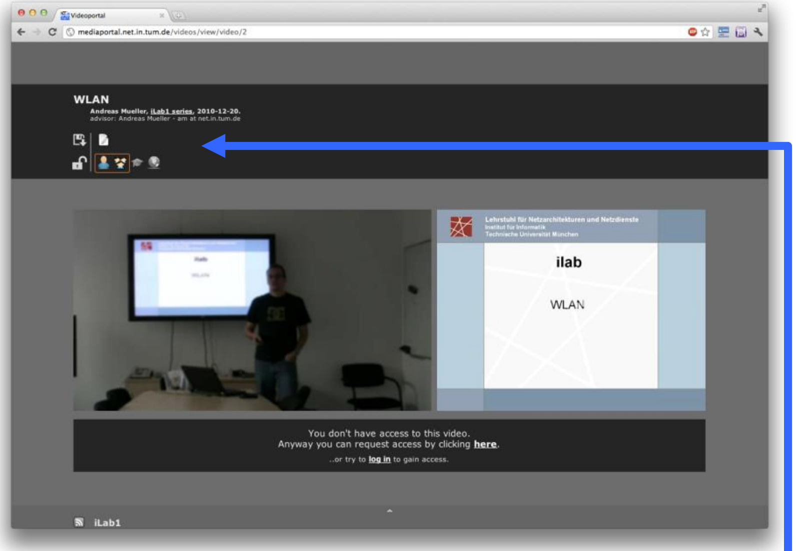
Videoportal
mediaportal.net.in.tum.de/videos/view/video/2
Lehrstuhl für Netzarchitekturen und Netzdienste
Institut für Informatik
Technische Universität München
ilab
WLAN
iLab1

WLAN

Andreas Mueller, iLab1 series, 2010-12-20.
advisor: Andreas Mueller - am at net.in.tum.de

You don't have access to this video.
Anyway you can request access by clicking **here**.
..or try to **log in** to gain access.

You fully control the access!
(Initially only you can access it!)

Grading – Influencing Factors

- First version of paper must be acceptable
  - Grade worse than 4.0: disqualification (seminar graded as 5.0)
- Observe the deadlines
  - Advisor meetings are compulsory
  - Use the upload form on Moodle course page for your submissions
  - 0.3 degrading per day for missed deadlines
- No submission
  - 1st version of paper: disqualification (Seminar graded as 5.0)
  - Other submissions: grade 5.0 for the concerning part
- Write the paper yourself
  - Plagiarism → disqualification (and we will check!)
  - Attempted cheating reported to the examination office
  - Summary when and why to cite: http://oxford.library.emory.edu/research-learning/citation-plagiarism/citing.html

FI vs IITM

Both seminars offer similar topics

- Again, check the proceedings (link on previous slide)
- The next FI Seminar will have an emphasis on advanced persistent threats

Future Internet

- Topic presentation: Friday, February 17, 14:00, room 03.07.023
  - This is just 2 days after the matching assignments are published!
- Seminar paper to be written in lecture free time (only few weeks until first version)
- Talks as block in the last week of the lecture free time
- Registration limited to 16 participants

Innovative Internet Technologies and Mobile Communications

- Topic presentation: Friday, April 28, 14:00, room 03.07.023
- Seminar paper to be written during the semester (until summer holidays)
- Talks weekly on Friday, 14:00, starting in June or July 2017 (probably 4 meetings)
- Registration limited to 14 participants

Seminar Proceedings

We publish your papers!

- If both you and your advisor agree
- Proceedings from the last years can be found on http://www.net.in.tum.de/publications/seminar%20proceedings/

Look at old proceedings

- Examples of papers we consider “good”
- Get an idea of the topics we cover

Best Paper Award

- We will choose a best paper in each seminar
- They will receive a certificate and a hardcopy of the proceedings

Registration

Registration is handled centrally on a dedicated web platform

1. You enter your seminar preferences into a dedicated web platform
2. We enter our student preferences
3. The system computes a **student-optimal** matching

- More info: http://docmatching.in.tum.de

If you want to be preferred by us put your matriculation number on the sheet handed out

- We will prefer you for both IITM and FI seminar
- The list is closed after this event!

The result of the matching is binding, i.e. **you cannot step down** from the course afterwards

- Only enter courses that you really want to participate in

Matching – Example 1

Your preferences

1. $your_favorite_seminar
2. $i8_seminar
3. $not_so_interesting_seminar

Putting your matriculation number on our list does **not** reduce your chance of being assigned to $your_favorite_seminar

It will however increase your chance of being assigned to $i8_seminar if you can’t be assigned to $your_favorite_seminar

Matching – Example 2

- Preferences Student A
  1. $very_popular_seminar
  2. $popular_seminar

- Preferences Student B
  1. $popular_seminar

Chances of getting assigned (to any course) are not higher for Student B compared to Student A

The Matching System works best **for you** if you honestly enter all your preferences

- Giving same priority to multiple courses is possible

About the Topics

All topics are derived from the ongoing research of our chair

- You may have a look at the research areas of our members: http://www.net.in.tum.de/members/

- Topics are often continuations of previous seminar topics
- Proceedings from the last years can be found on http://www.net.in.tum.de/publications/seminar%20proceedings/

Topics will be assigned according to your preferences similar to the faculty-wide process of seminar matching

- Each topic is unique and moved to the next seminar only if not selected previously
- We aim for +20% topics compared to seats in the seminar

- The following slides show topics that will be offered in the next seminar (further topics will be announced)

Focus Topic: Advanced Persistent Threats (APT)

Examples:

- Stuxnet, Operation Shady RAT, Operation Aurora

What makes an attack to an APT? Differentiation:

- Target a specific entity (company, person)
- Have a high degree of covertness over a long period of time
- Sophisticated techniques used to compromize systems

Some typical stages

- Initial compromise (spear phishing, waterhole attack, zero day)
- Establish foothold (backdoor, tunnels)
- Escalate privileges (exploits, password cracking)
- Internal network reconnaissance
- Laterally movement (attack is spreading)
- Complete mission (exiltrate data, cover tracks)

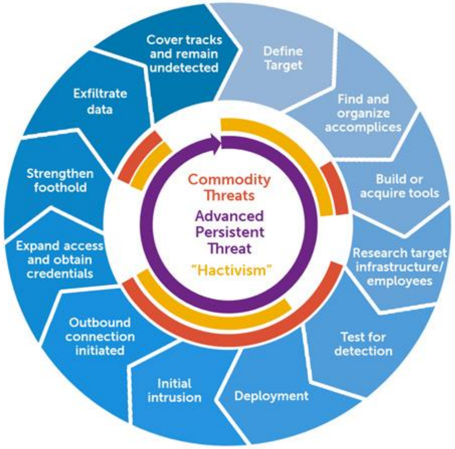
Cover tracks and remain undetected
Define Target
Find and organize accomplices
Build or acquire tools
Research target infrastructure/ employees
Test for detection
Deployment
Initial intrusion
Outbound connection initiated
Expand access and obtain credentials
Strengthen foothold
Exfiltrate data
Commodity Threats
Advanced Persistent Threat
"Hactivism"

Focus Topic: Advanced Persistent Threats (APT)

Examples:

- Stuxnet, Operation Shady RAT, Operation A

What makes a

- Target a spe
- Have a high c
- Sophisticated

Some typical stag

- Initial compromi ay)
- Establish foothol)
- Escalate privileges (exploits, password cracking)
- Internal network reconnaissance
- Laterally movement (attack is spreading)
- Complete mission (exiltrate data, cover tracks)

Offered in the blockseminar only!

We will issue BA/MA theses in this area soon

Your opportunity to familiarize yourself with a highly relevant topic!

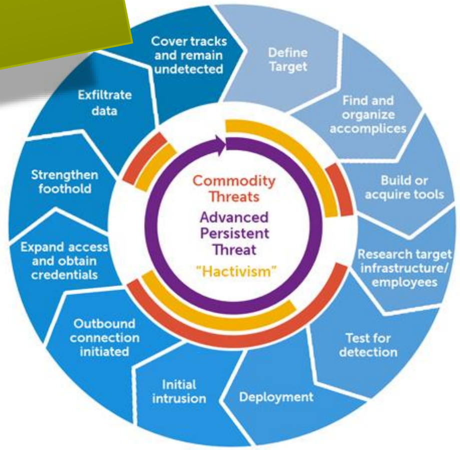
Cover tracks and remain undetected
Define Target
Find and organize accomplices
Build or acquire tools
Research target infrastructure/ employees
Test for detection
Deployment
Initial intrusion
Outbound connection initiated
Expand access and obtain credentials
Strengthen foothold
Exfiltrate data
Commodity Threats
Advanced Persistent Threat
"Hactivism"

seminar@net.in.tum.de

Industrial Espionage & APT Malware Zoo (Johannes)

Thyssenkrupp, EADS; GhostNet, Careto, Stuxnet, Duqu, Flame, …
Targeted Attacks, primary Goals: Information Extraction (i.e. not DDoS, Cryptomalware, Adware, etc.)

Your Task:

- Provide an overview over malware for information extraction
- Document known, high impact malware campaigns in recent time (last 10 years)
- What is known about the malware, techniques and targets?
- How was the malware ultimately discovered?

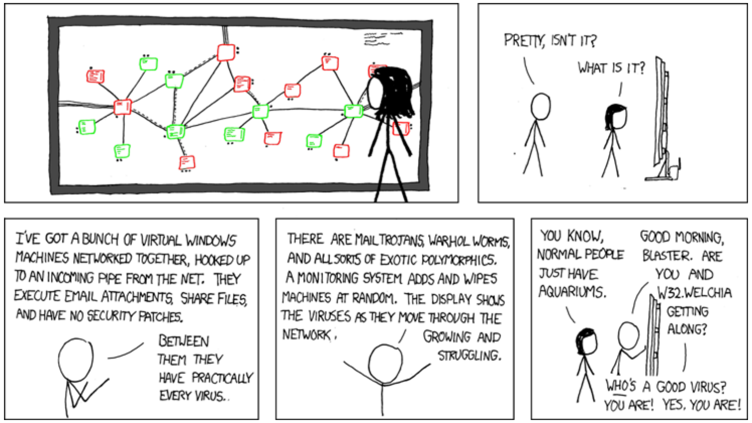
PRETTY, ISN'T IT?
WHAT IS IT?
I'VE GOT A BUNCH OF VIRTUAL WINDOWS MACHINES NETWORKED TOGETHER, HOOKED UP TO AN INCOMING PIPE FROM THE NET. THEY EXECUTE EMAIL ATTACHMENTS, SHARE FILES, AND HAVE NO SECURITY PATCHES.
BETWEEN THEM THEY HAVE PRACTICALLY EVERY VIRUS.
THERE ARE MAIL TROJANS, WARHOL WORMS, AND ALL SORTS OF EXOTIC POLYMORPHICS. A MONITORING SYSTEM ADDS AND WIPES MACHINES AT RANDOM. THE DISPLAY SHOWS THE VIRUSES AS THEY MOVE THROUGH THE NETWORK,
GROWING AND STRUGGLING.
YOU KNOW, NORMAL PEOPLE JUST HAVE AQUARIUMS.
GOOD MORNING, BLASTER. ARE YOU AND W32.WELCHIA GETTING ALONG?
WHO'S A GOOD VIRUS? YOU ARE! YES, YOU ARE!

Sources: Blogs of Symantec, F-Secure, Kaspersy, *-CERT reports etc.
Check for scientific sources

Protect yourself: Use of gpg, tor required.

Image: https://www.xkcd.com/350/

Sophisticated Approaches to Penetrate Highly Secured Systems (Liebald)

Modern systems are usually protected against common malware threats.

However, malware can apply different mechanisms to still infiltrate such a system and avoid detection.

```
a.fn.scrollspy=d,this},a(window).on(
y),+function(a){"use strict";function b(b){return this.each(function(){var
e[b]()})}var c=function(b){this.element=a(b)};c.VERSION="3.3.7",c.TRANSITION_DURATION=150,c.prot
opdown-menu)"),d=b.data("target");if(d||(d=b.attr("href"),d=d&&d.replace(/.*(?=#[^\s]*$)/,"")),!
st a"),f=a.Event("hide.bs.tab",{relatedTarget:b[0]}),g=a.Event("show.bs.tab",{relatedTarget:e[0]
faultPrevented()){var h=a(d);this.activate(b.closest("li"),c),this.activate(h,h.parent(),functio
rigger({type:"shown.bs.tab",relatedTarget:e[0]})})}}},c.prototype.activate=function(b,d,e){func
> .active").removeClass("active").end().find('[data-toggle="tab"]').attr("aria-expanded",!1),
ia-expanded",!0),h?(b[0].offsetWidth,b.addClass("in")):b.removeClass("fade"),b.parent(".dropdow
().find('[data-toggle="tab"]').attr("aria-expanded",!0),e&&e()}var g=d.find("> .active"),h=e&&
e")||!!d.find("> .fade").length);g.length&&h?g.one("bsTransitionEnd",f).emulateTransitionEnd
;var d=a.fn.tab;a.fn.tab=b,a.fn.tab.Constructor=c,a.fn.tab.noConflict=function(){return a.fn.t
"show")};a(document).on("click.bs.tab.data-api",'[data-toggle="tab"]',e).on("click.bs.tab.data
se strict";function b(b){return this.each(function(){var d=a(this),e=d.data("bs.affix"),f="ob
typeof b&&e[b]()})}var c=function(b,d){this.options=a.extend({},c.DEFAULTS,d),this.$target=a
",a.proxy(this.checkPosition,this)).on("click.bs.affix.data-api",a.proxy(this.checkPositionWi
ull,this.pinnedOffset=null,this.checkPosition()};c.VERSION="3.3.7",c.RESET="affix affix-top
State=function(a,b,c,d){var e=this.$target.scrollTop(),f=this.$element.offset(),g=this.$targ
"bottom"==this.affixed)return null!=c?!(e+this.unpin<=f.top)&&"bottom":!(e+g<=a-d)&&"bottom"
!=c&&e<=c?"top":null!=d&&i+j>=a-d&&"bottom"},c.prototype.getPinnedOffset=function(){if(this
.RESET).addClass("affix");var a=this.$target.scrollTop(),b=this.$element.offset();return
WithEventLoop=function(){setTimeout(a.proxy(this.checkPosition,this),1)}
nt.height(),d=this.options.offset,e=d.top,f=d.botto
eof e&&(e=d.top(this.$element)),"f
nt.css("top",""
```

https://static.pexels.com/photos/239898/pexels-photo-239898.jpeg

Your Task:

- What are typical and advanced malware defense mechanisms?
- On which properties of malware do they rely on?
- How can malware penetrate such secured systems and avoid detection for a longer period of time?

Keywords:

- Malware obfuscation, signature detection, social engineering [1], oligomorphic/polymorphic/metamorphic code [2], …

[1] Sherly Abraham, InduShobha Chengalur-Smith, An overview of social engineering malware: Trends, tactics, and implications, Technology in Society, August 2010

[2] I. You and K. Yim, "Malware Obfuscation Techniques: A Brief Survey," 2010 International Conference on Broadband, Wireless Computing, Communication and Applications, Fukuoka, 2010

SotA Analysis of APT Defense [Kinkelin]

- APTs are serious problems for companies and organizations
- Hence, the field of APT defense is a trending topic for academia and a hot market for product vendors
- Your task is a **state of the art analysis** that collects **APT detection approaches** described in **academic papers** and on **product web pages**

- Examples *include*
  - file/packet-based approaches [1]
  - context-based approaches [2]
  - flow-based approaches [3]

- Your goals:
  - Explain why existing cyber attack defense is not able to defend against APTs
  - Create a taxonomy of approaches (e.g. methods, effort, data privacy, availiablity, ...)
  - Showcase selected academic/commercial approaches

[1] Palo Alto Wildfire: https://www.paloaltonetworks.com/products/secure-the-network/subscriptions/wildfire

[2] Giura, P., & Wang, W: A context-based detection framework for advanced persistent threats.

[3] Marchetti, M et al: Analysis of high volumes of network traffic for Advanced Persistent Threat detection

Pictures: https://commons.wikimedia.org/wiki/GNOME_Desktop_icons

Modelling Advanced Persistent Threats (Oliver)

Advanced Persistent Threats

- APTs are specifically tailored compromise attempts
- Focus on well-defined target
- Frequently involves exploiting multiple 0days
- Examples: Stuxnet [1], Duqu [2]
- Modelling APTs can advance detection mechanisms

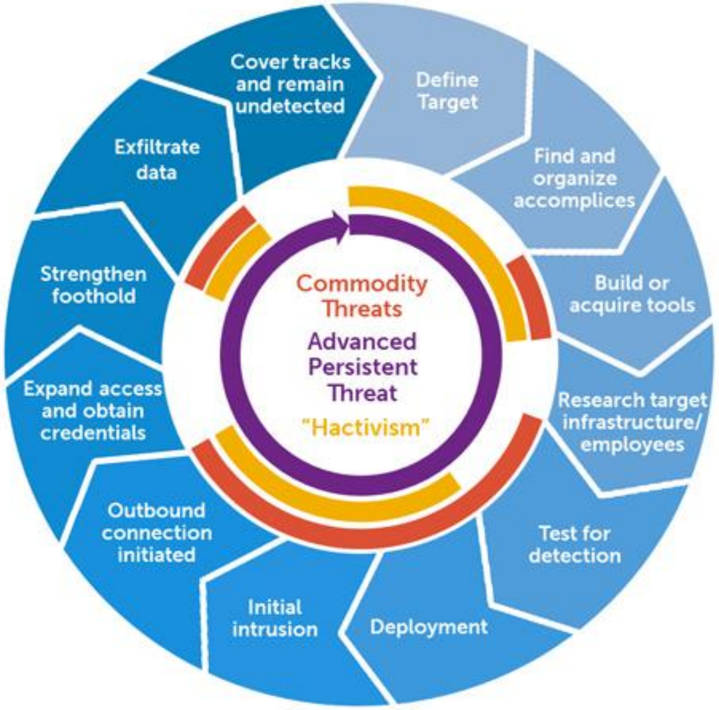
Cover tracks and remain undetected
Define Target
Find and organize accomplices
Build or acquire tools
Research target infrastructure/ employees
Test for detection
Deployment
Initial intrusion
Outbound connection initiated
Expand access and obtain credentials
Strengthen foothold
Exfiltrate data
Commodity Threats
Advanced Persistent Threat
"Hactivism"

Your Task:

- Get familiar with APTs
- Research modelling languages (e.g. STIX [3]) and identify suitable one
- Model APT using chosen modelling language
- Identify possible feature vectors derived from model to aid APT detection

[1] S. Karnouskos: “Stuxnet worm impact on industrial cyber-physical system security”, IECON’11.

[2] B. Bencsáth: “Duqu: Analysis, detection, and lessons learned”, EuroSec’12.

[3] S. Barnum: “Standardizing cyber threat intelligence information with the Structured Threat Information eXpression (STIX™)”

Spear Phishing and Watering Holes (Heiko)

- Initial way to compromise a victim network
- Spam Phishing vs Spear Phishing
- Watering Hole Attack

Your Task:

- Concept of Phishing in the context of persistent threats, related concept of watering hole attacks
- How do Spear Phishing mails look like?
- Who is targeted?
- What information is used?
- How does the system get compromised?
- What can be done against it?

Strategies for Malware in Cyber Conflicts (Heiko)

Persistent Threats and Cyber Conflict

- Weaponry used
- Zero Day Exploits
- Persistence
- Stealth
- Victory and Defeat
- Targeted Attack

Your Task:

- What is a cyber conflict?
- How does it compare to non-cyber conflicts?
- What tools are used?
- Strategies to consider when using the tools
- Evidence / how to detect
- Advanced Persistent Threats – how are they related to the issue.

Security of CPE Management Protocols (Oliver)

Remote CPE WAN Management Protocols

- Configuration of CPE (e.g. your home modem/router)
- Firmware upgrade
- Diagnostics and troubleshooting
- Example: TR-069

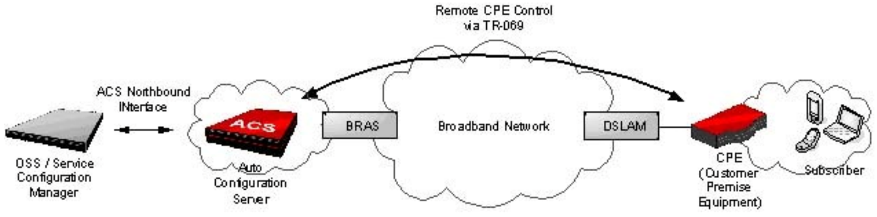

Your Task:

- Get familiar with remote CPE management approaches
- Find and compare different protocols from a security perspective
- What is their threat model?
- How do these protocol achieve the security guarantees?
- What went wrong during successful attacks [1]?

[1] http://www.linus-neumann.de/2016/11/30/warum-die-telekom-router-ausgefallen-sind/

Penetrating secured systems (Liebald)

Modern systems are usually protected against common malware threats.
However, there are mechanisms for malware to avoid detection.

Your Task:

- Get an overview about how computer systems can be attacked by malware
- What are typical defense mechanisms?
- On which properties of malware do they rely on?
- How can malware penetrate secured systems and avoid detection for a longer period of time?

https://static.pexels.com/photos/239898/pexels-photo-239898.jpeg

Keywords:

- Malware obfuscation, signature detection, social engineering [1], oligomorphic/polymorphic/metamorphic code [2], …

[1] Sherly Abraham, InduShobha Chengalur-Smith, An overview of social engineering malware: Trends, tactics, and implications, Technology in Society, August 2010

[2] I. You and K. Yim, "Malware Obfuscation Techniques: A Brief Survey," 2010 International Conference on Broadband, Wireless Computing, Communication and Applications, Fukuoka, 2010

Location Privacy Preserving Mechanisms (Totakura)

Location based services have proven to be useful in our day-to-day life. However, since our location data is exposed to the service, they raise concerns over user privacy. Location Privacy Preserving Mechanisms (LLPM) try to protect our privacy while still being able to provide the service. Some of them are:

- Spatial cloaking[1]
- Sending dummy queries [2]
- Mix-zones [3]

Your tasks:

- Get familiar with LLPM
- Survey literature for protection mechanisms
- Survey literature for attacks agains LLPM
- Compare LLPMs

Starting Points:

[1] Gruteser, Marco, and Dirk Grunwald. *"Anonymous usage of location-based services through spatial and temporal cloaking."* Proceedings of the 1st international conference on Mobile systems, applications and services. ACM, 2003.

[2] Shankar, Pravin, Vinod Ganapathy, and Liviu Iftode. "Privately querying location-based services with SybilQuery." Proceedings of the 11th international conference on Ubiquitous computing. ACM, 2009.

[3] Beresford, Alastair R., and Frank Stajano. "Location privacy in pervasive computing." IEEE Pervasive computing 2.1 (2003): 46-55.

Passive Host Fingerprinting (Scheitle, Diekmann)

Like nmap -O , but passive

Attacker model: a router, sniffing the network

- Why passive?
- How uniquely can we identify hosts?
- How to recognize hosts previously seen?
- What features are available? What works well?

Your Task:

- Introduce and compare active and passive host fingerprinting (1 page)
- Find and compare different approaches for passive host fingerprinting in literature (3pp.)
- Distill and compare kinds of data related work relies on? (2pp.)
- What limitations do existing approaches have? Can you think of ways for improvement or new approaches? (1p.)

- Related Work: nmap (active), https://arxiv.org/abs/1610.07251, https://arxiv.org/abs/1606.07613,

Post-Action Verification in Self-Organizing Networks

Tsvetkov

Problems of mobile SONs

- Increasing reliance on SON features to perform the correct optimization tasks
- The impact of each function’s action on the environment depends upon the actions of other functions as well

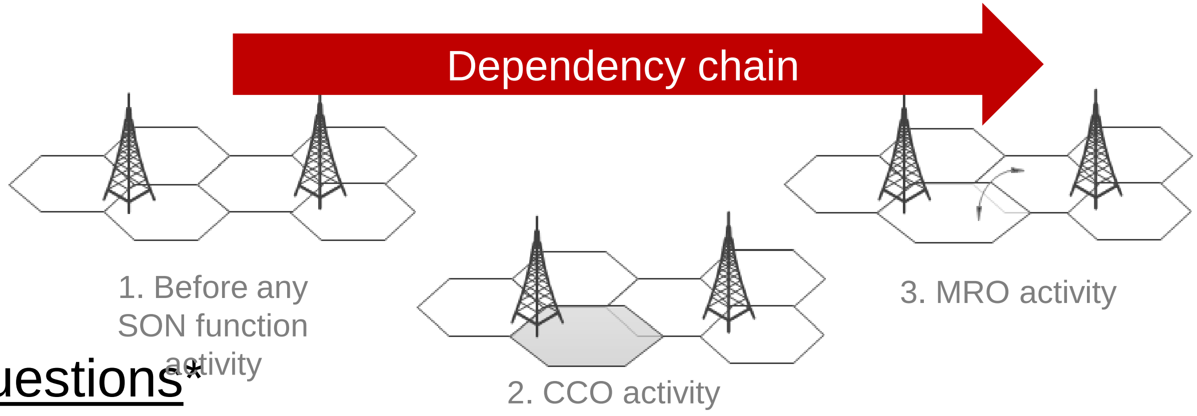

Tasks & Questions*

- Existing SON verification solutions?
- Used anomaly detection techniques?
- Used diagnosis approaches?
- What about corrective actions?

Consequences in case of a wrong decision?

(*) Writing in English is preferable.

Semantic integration of IoT devices (Jan)

IoT devices use a multitude of protocols. Integrating such devices is difficult to do automatically. By using semantic web concepts:

- Protocols can be translated automatically [1]
- Data can be discovered with complex queries
- Complex services can be orchestrated automatically.

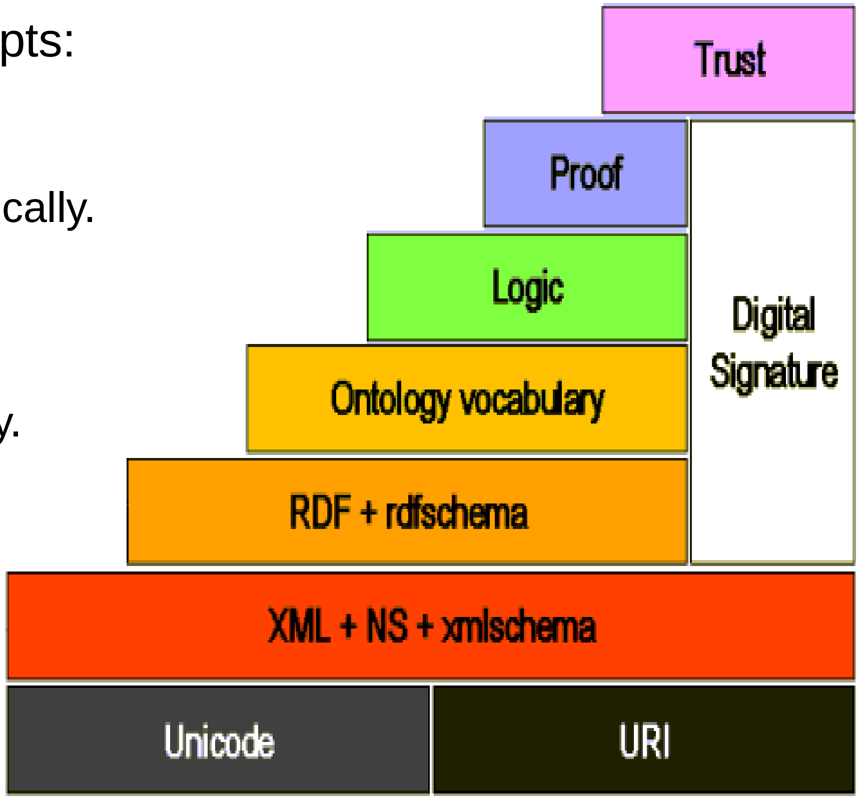

Your Task:

- Familiarize yourself current semantic technology.
- Find and summarize semantic IoT integration
  approaches in literature.
- What limitations are there?
  (performance, expressivity, etc.)
- Semantically model a simple device interaction.
  (i.e. light switch to lamp)

[1] D. O'Sullivan and D. Lewis, 'Semantically Driven Service Interoperability for Pervasive Computing', in *Proceedings of the 3rd ACM International Workshop on Data Engineering for Wireless and Mobile Access*, New York, NY, USA, 2003, pp. 17–24.

Caching the Internet (Liebald)

Caching plays an important, yet usually invisible role in the Internet:

- Reduced delay
- Decreased server load
- Enhanced Reliability/Resilience

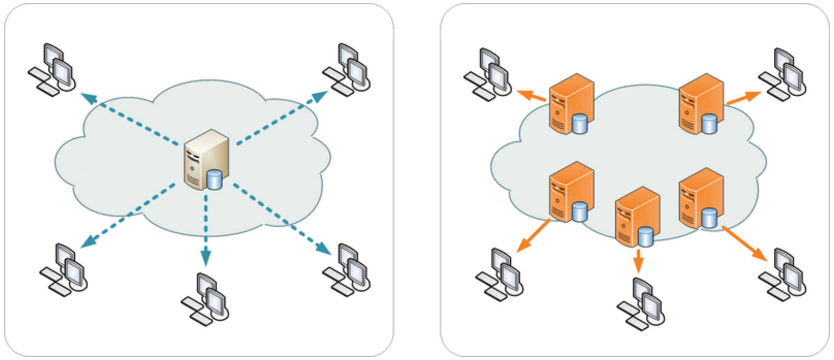

https://en.wikipedia.org/wiki/Content_delivery_network#/media/File:NCDN_-_CDN.png

Your Task:

- Get familiar with the concept
- of caching in the Internet
- Find and compare different architectures which utilize caching in literature
  - CDN, ICN, P2P,…
- What kind of mechanisms do they rely on?
- What are benefits of different approaches?
- What drawbacks can caching introduce?

Detecting load balancers on the Internet (English only)

Rouhi, Scholz

Previous works [1, 2, 3] have demonstrated that

- Proxies and load balancers are commonplace in today’s Internet
- Widely deployed for popular content providers

Your Tasks:

- Get familiar with the concept
- What types of techniques are used for load balancing?
- On what layers [4] do the techniques operate?
- Compare different approaches of detecting load balancers
- Point out Pros and Cons of their deployment and possible remedies

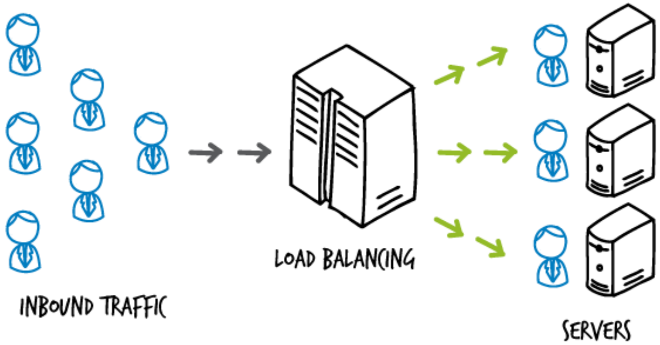

http://tinyurl.com/zeabqau

[1] http://ieeexplore.ieee.org.eaccess.ub.tum.de/stamp/stamp.jsp?arnumber=4579532
[2] http://tinyurl.com/h53auly
[3] https://www.sans.org/reading-room/whitepapers/testing/identifying-load-balancers-penetration-testing-33313
[4] https://github.com/jmbr/halberd

SCION Internet Architecture (Johannes)

SCION (Scalability, Control, and Isolation on Next-Generation Networks) is an architecture developed at the ETH Zürich as a clean-slate internet design.
Main Goals: provide a secure and resilient “internet” with a minimal TCB.

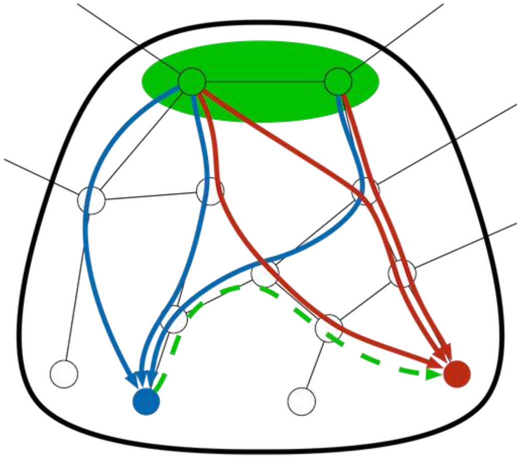

Your Task:

- What is the project about?
- What are it’s goal?
- What is actually achieved, does the code run?
- How does it compare to the current internet?
- What problems are solved?

Requirements: Understanding how the internet works in the default free zone, i.e. you can tell me something about AS, BGP, peering, transit.

Starting Point: http://scion-architecture.net/

Image: https://www.scion-architecture.net/img/domain-logo.png

Open Source Routers and New Network Protocols (Cordeiro)

Open source routers largely use by small and medium autonomous systems and as route server by internet exchange points.

Also very interesting for research purposes, as computers and virtual machines can be used, instead of expensive routers with inflexible proprietary software.

Unfortunately, as common in various open source projects, it lacks documentation.

Your Task:

- Get familiar with Quagga router architecture
- Implement a skeleton module (C code) for a new routing protocol for Quagga
- Document the steps for creating this basic new module in a tutorial style

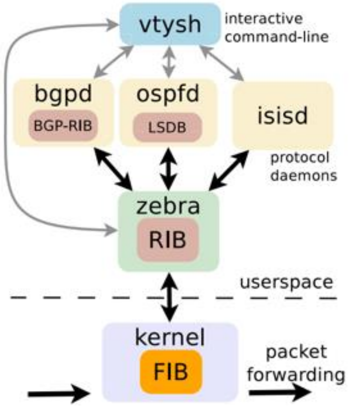

[1] Quagga Routing Suite - http://www.nongnu.org/quagga/docs.html

[2] ZEBRA FOR DUMMIES - http://www.nongnu.org/quagga/zhh.html - last update: Feb. 2011

[3] P. Jakma, D. Lamparter; Introduction to the Quagga Routing Suite; 2014; https://www.academia.edu/5754254/Introduction_to_the_Quagga_Routing_Suite

Talk and Paper in English is preferable

Resource management with learning algorithms

(Cora)

Recent works [1] have shown that learning algorithms can be used to

- Decentralise network resource management
- Dynamically and opportunistically allocate resources depending on demand
- Improve maximum number of network requests

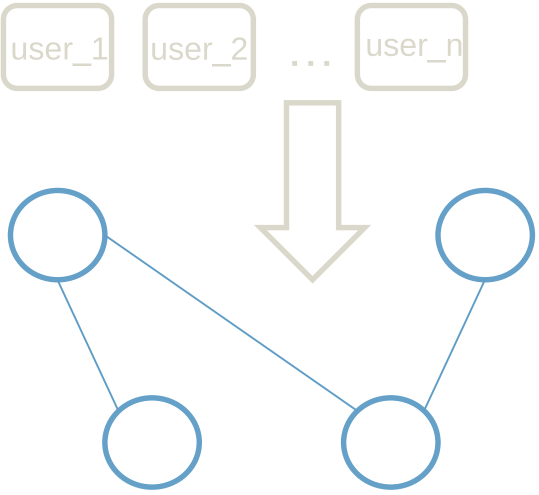

Your Tasks:

- Get familiar with network resource management and Q-learning
- Find related work (e.g. using different algorithms). How does it compare to the results presented in [1]?
- What are the advantages of using learning algorithms compared to traditional methods?
- What limitations are there to the proposed approach? How could these be remedied?

[1] R. Mijumbi, J. L. Gorricho, J. Serrat, M. Claeys, F. De Turck and S. Latré, "Design and evaluation of learning algorithms for dynamic resource management in virtual networks," *2014 IEEE Network Operations and Management Symposium (NOMS)*, Krakow, 2014, pp. 1-9.

Comparison of Scientific Workflow Systems (Wohlfart, Raumer)

Scientific workflow systems structure and automate data analysis as a series of operations applied to the data

Benefits:

- Documented, reproducible process
- Simplified sharing/reuse between scientists
- Hide complexity

We operate a networking testbed

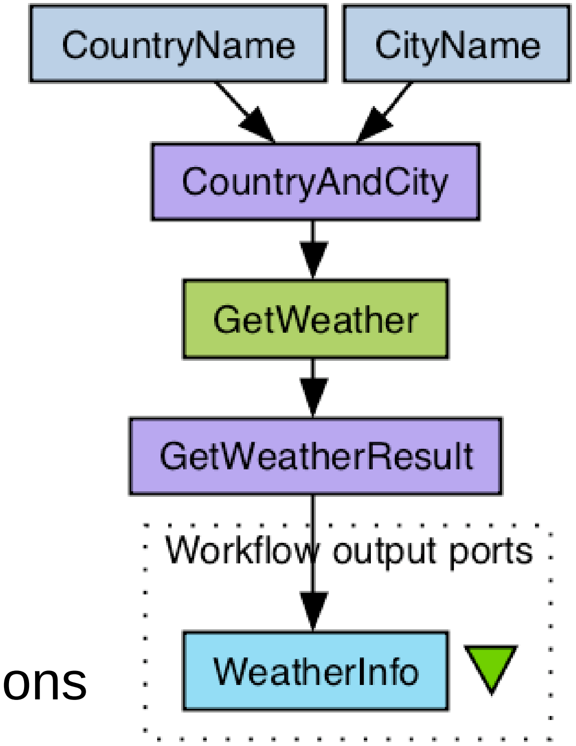

Source: http://www.taverna.org.uk/introduction/why-use-workflows/

Your Task:

- Understand the benefits and problems of SWS in our testbed
- Analyze and compare selected SWS implementations
  - e.g. Apache Taverna, Kepler, VisTrails, LONI Pipeline

[1] V. Curcin et.al "Scientific workflow systems - can one size fit all?". 2008 Cairo Int. Biomedical Engineering Conference

Baltinet: Mininet on real hardware (Wohlfart, Raumer)

“Mininet creates a realistic virtual network, running real kernel, switch and application code, on a single machine” [1,2]

- Python script defines topology and commands executed on the hosts
- Baltinet is a Mininet like approach for experiments in the Baltikum Testbed [3]

Your Task:

- Sell the idea of Mininet
- Case study: apply Mininet experiment to the Baltikum testbed

[1] B.Lantz, B.Heller, N.McKeown “A Network in a Laptop: Rapid Prototyping for Software-Defined Networks”, ACM HotNets, 2010
[2] http://mininet.org/
[3] C.Feiler, “Das Baltikum Testbed,” FI/IITM seminar, vol. NET-2016-09-1, 2016

Queueing Theory Models for Performance of X86 Interconnect Devices

(Daniel)

Queueing theory is the mathematical study of waiting lines, or queues. In queueing theory a model is constructed so that queue lengths and waiting time can be predicted.

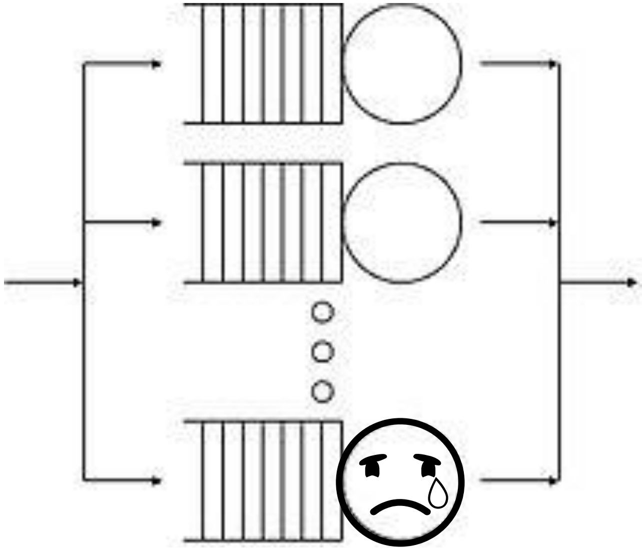

Your Task:

- Survey literature
- Introduce queuing theory
- Describe how systems are mapped to models.
- Focus on the performance aspects of a single server

Starting points:

[1] R.Bruschi, F.R.Davoli, P.Lago, J.F.Pajo “*Joint Power Scaling of Processing Resources and Consolidation of Virtual Network Functions*” CloudNet, 2016
[2] T.Meyer, F.Wohlfart, D.Raumer, B.E.Wolfinger, G.Carle, “*Validated Model-Based Prediction of Multi-Core Software Router Performance*”, PIK, 2014
[3] S.Gebert, T.Zinner, S.Lange, C.Schwartz, P.Tran-Gia“Performance Modeling of Softwarized Network Functions Using Discrete-Time Analysis. ITC, 2016

Using Service Curves To Describe Software-based Network Functions

(Daniel, Heiko)

Network Calculus (NC) describes systems and events with a cumulative function of the time. In NC-theory the term Service Curves refers to the interval bound function for the available resources of a system.

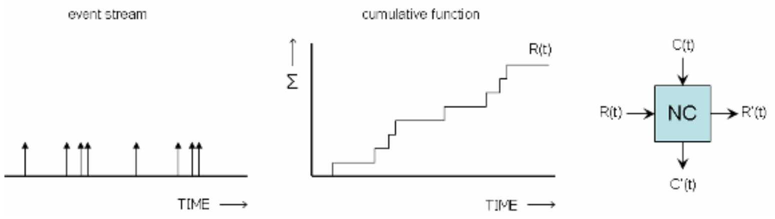

Your Task:

- Survey literature
- Introduce Network Calculus
- Describe how deployed network functions are described with Service Curves
- How are model parameters measured on real systems?
- *Bonus: how can they be applied to predict performance of network function chains?*

Starting point:

[1] A.Undheim, Y.Jiang, P.J.Emstad *“Network Calculus Approach to Router Modeling with External Measurements”* VALUETOOLS 2011

Increasing VM networking performance with virtio

(Paul, Daniel)

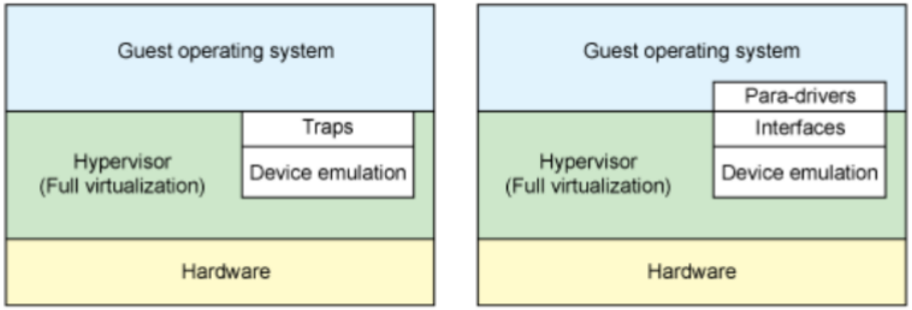

Source: M. Tim Jones, "Virtio: An I/O virtualization framework for Linux Paravirtualized I/O with KVM and lguest", 2010

Background/Interest in OS/network driver helpful advised

Your Task:

- What virtio?
- Why virtio?
- How virtio?
- additionally: older/full virtualization and newer/scientific alternatives

Starting Point: Google Deutschland

Effects of Unaligned Memory Access on the Performance of Packet Analyzers (Paul)

Packets have odd sizes and you need to store a large number of them in memory… what do you do?

- Align them nicely on x-byte boundaries?
- Align some protocol header on some boundary?
- Just Store them back-to-back

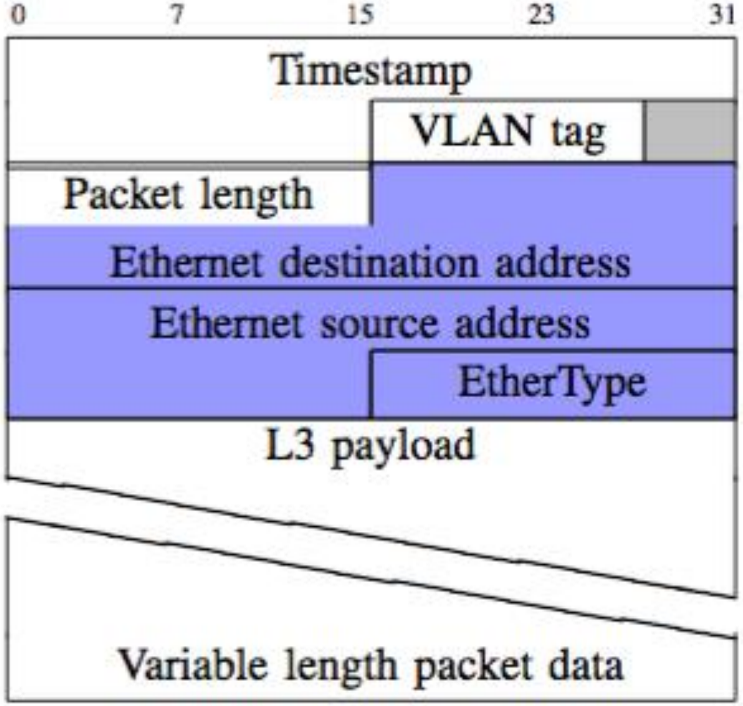

Your Task:

- Extend an existing benchmark with more scenarios
- Test different scenarios on different CPUs
- Read CPU performance counters to find bottlenecks
- Research background on unaligned memory access and alignment in popular packet storage formats

Necessary skills

- Experience with in-memory data analysis
- Experience with profiling
- Experience with CPU architecture

I‘m following my happy little packet through the Linux kernel until it gets dropped… (Raumer, Schwaighofer*)

Describe how packets are processed by Linux and where firewall functionality is added. Focus on kernel versions newer then 3.13 and changes since that version. Explain nftables, BPF!

Do not take that topic if you are afraid of ugly code!

*) not responsible for the title

Further topics will be presented in time