Karlsruhe Institute of Technology

Resilient Networking

Disclaimer: this course has been created with very valuable input from Günter Schäfer, Mathias Fischer, Michael Rossberg, and the members of the Chair

Module 1 – Preliminaries (Winter Term 2022)

Thorsten Strufe

Competence Center for Applied Security Technology

Lecture Outline

- Who are we?
- Organizational matters (preliminaries)
- Course outline

- A brief introduction

Who is Who

- Chair of „Privacy and IT Security“
- For this lecture:
- Thorsten Strufe (Lectures)
  - 50.34/281
  - thorsten.strufe [at] kit.edu

- Teaching assistants
- This lecture doesn‘t have one.

- Consultation:
- Send me an email (repeatedly…)

- https://ps.kastel.kit.edu/

Karlsruhe Institute of Technology

KASTEL

What motivates us at the chair…?

Humanity and Cultural Practices

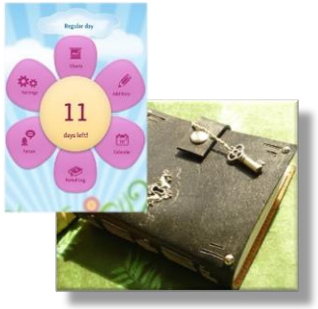

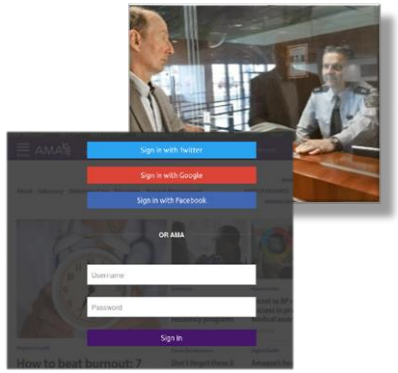

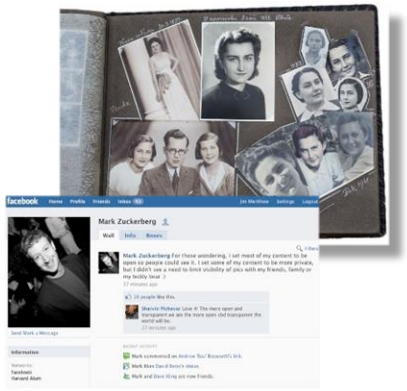

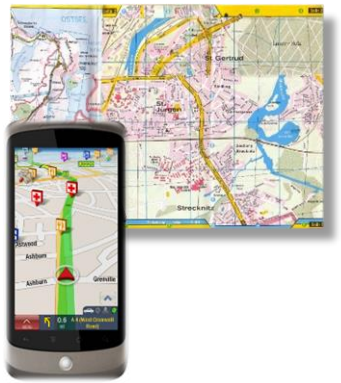

#TactileInternet

ceti.one

Access: Type, Scope, and Trust

1: Personal, unidentified
2: Local, decentralized
3: Trust in direct peer (village)

Access: Type, Scope, and Trust Today

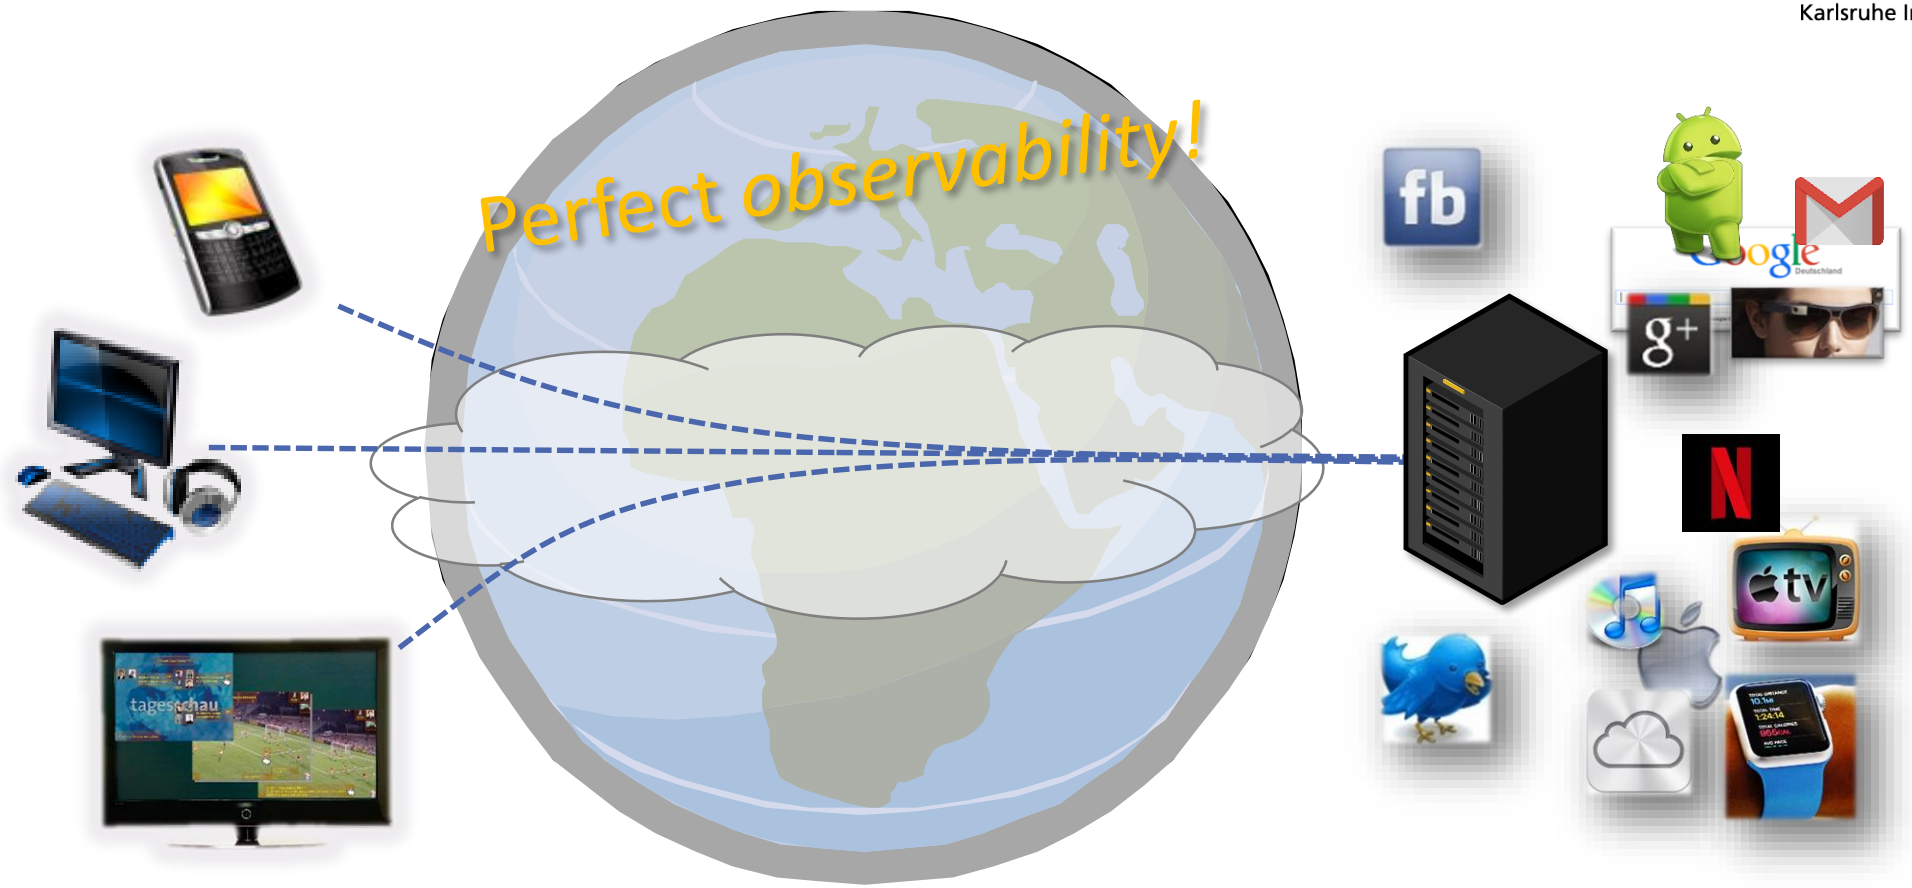

1: Central, unique global login services
2: Global access over Internet
3: Trust in … (I)SP?

Modelling System and Adversaries

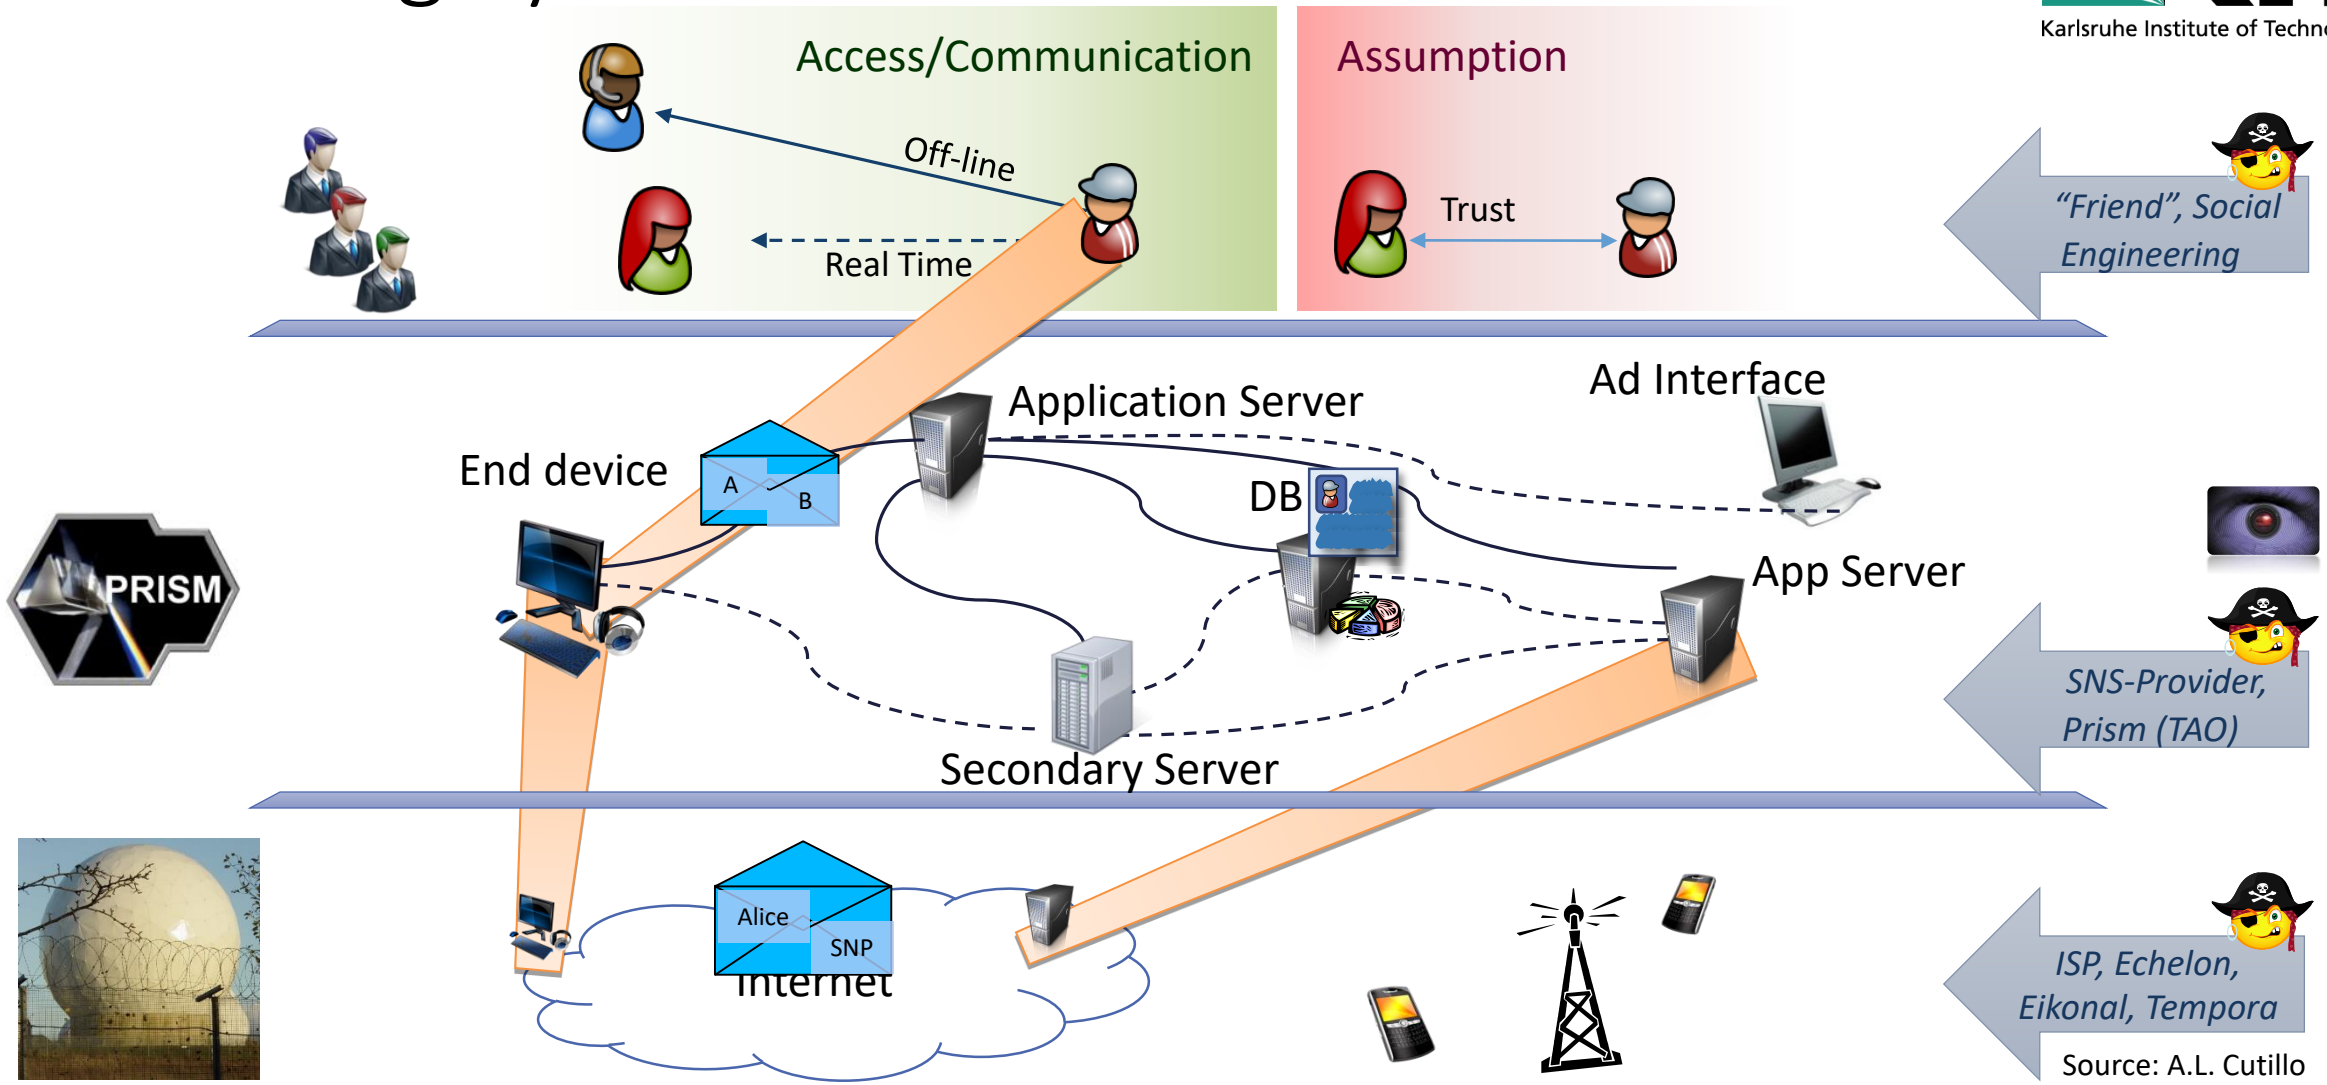

What we're working on…

- ***User understanding***

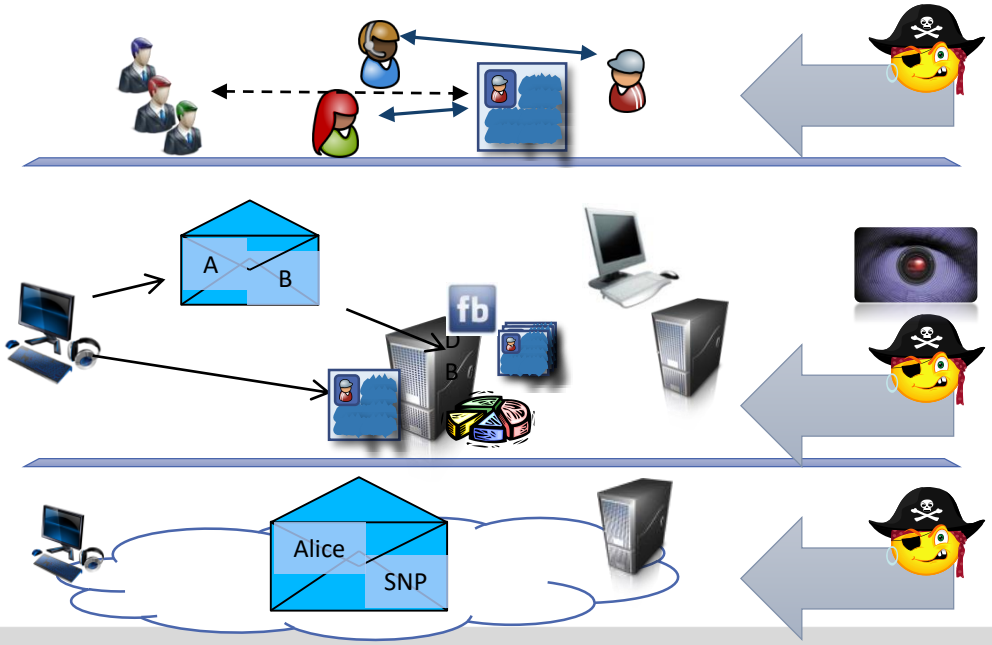

[7] FC '11
[8] WWW '12

What we're working on…

- ***User understanding***
  - Intention recognition
  - Privacy analyses
  - Data sanitation

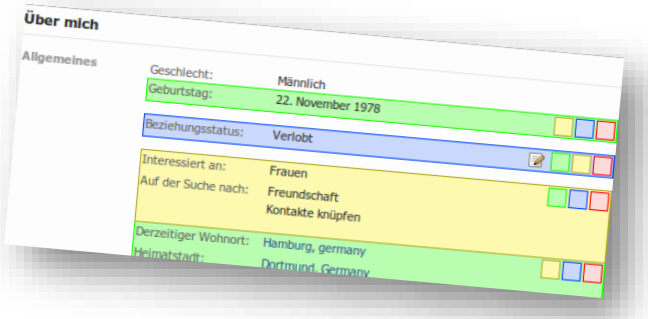

[redacted]

[redacted]

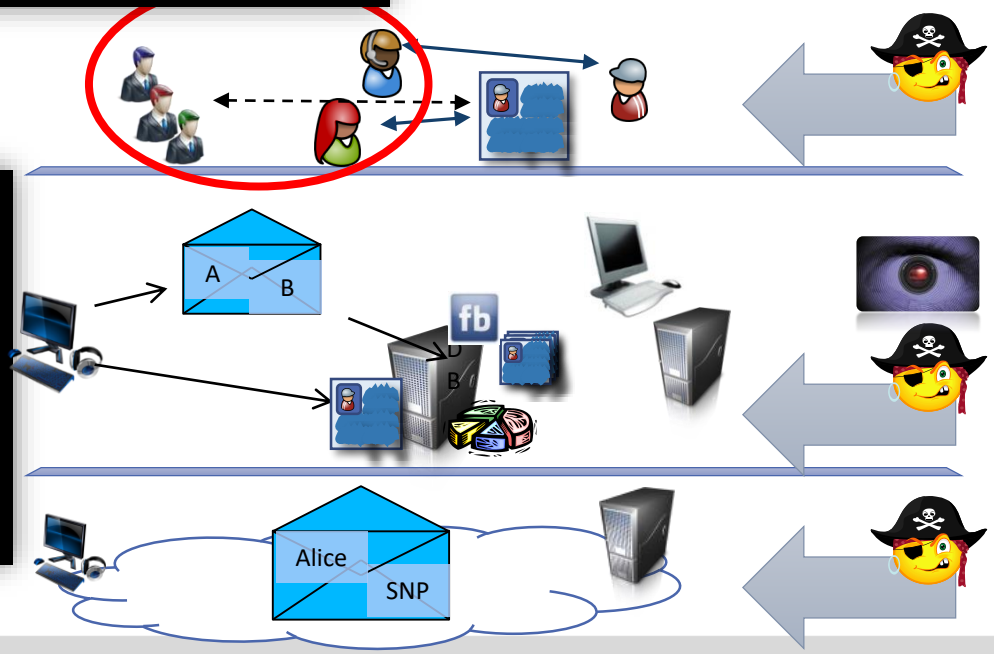

[7] FC '11
[8] WWW '12

Identifiability on the Web

- Web-Tracking is ubiquitous

- Situation:
  - Tracker claim anonymity
  - „delete last octett“: generalization
  - GDPR: Pseudonym ≠ Anonym

- Study
  - Cooperation with private partner
  - Comprehensive data set (German Web, 2-3 Bn visits per day)
  - Questions:
    - To which extent is behavior a pseudonym?
    - How little is needed to identify a trace?

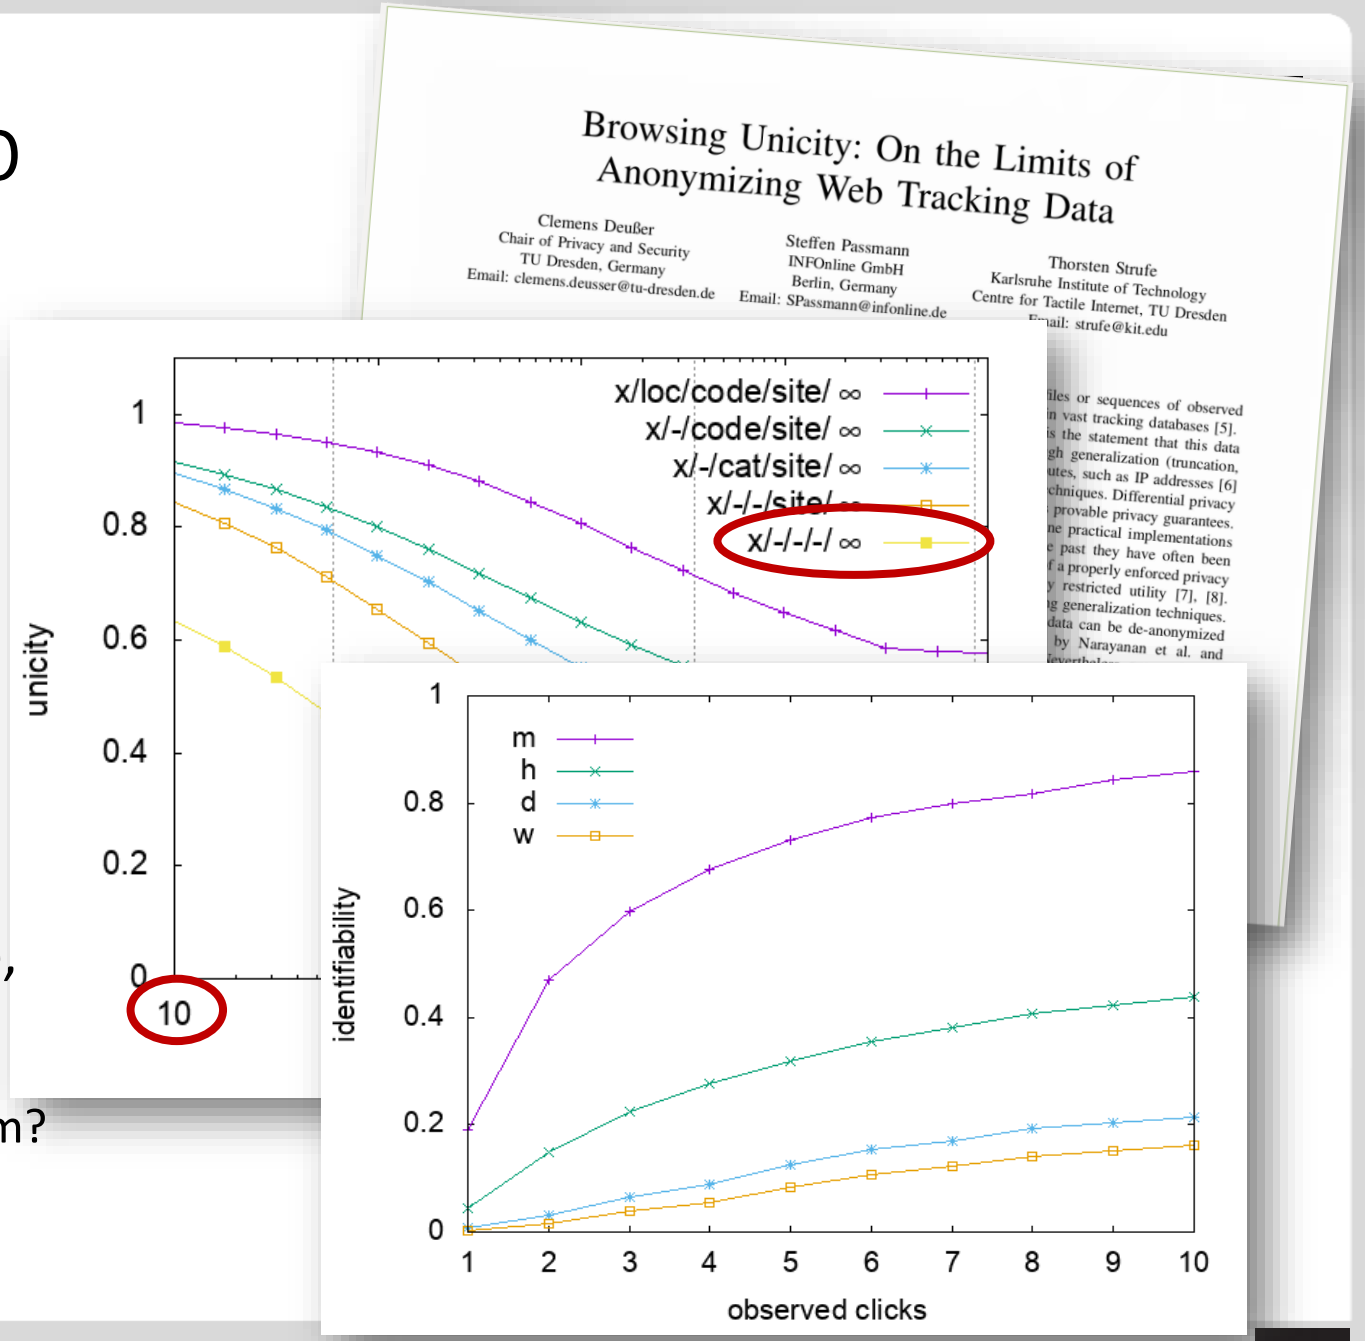

[9] IEEE S&P ’20

What we're working on…

- ***User understanding***
  - Intention recognition
  - Privacy analyses
  - Data sanitation

- ***Privacy-Enhancing Technologies***
  - Anonymity metrics
  - Anonymous services (f2f/Web)
  - Anonymous Communication (Tor, ..)

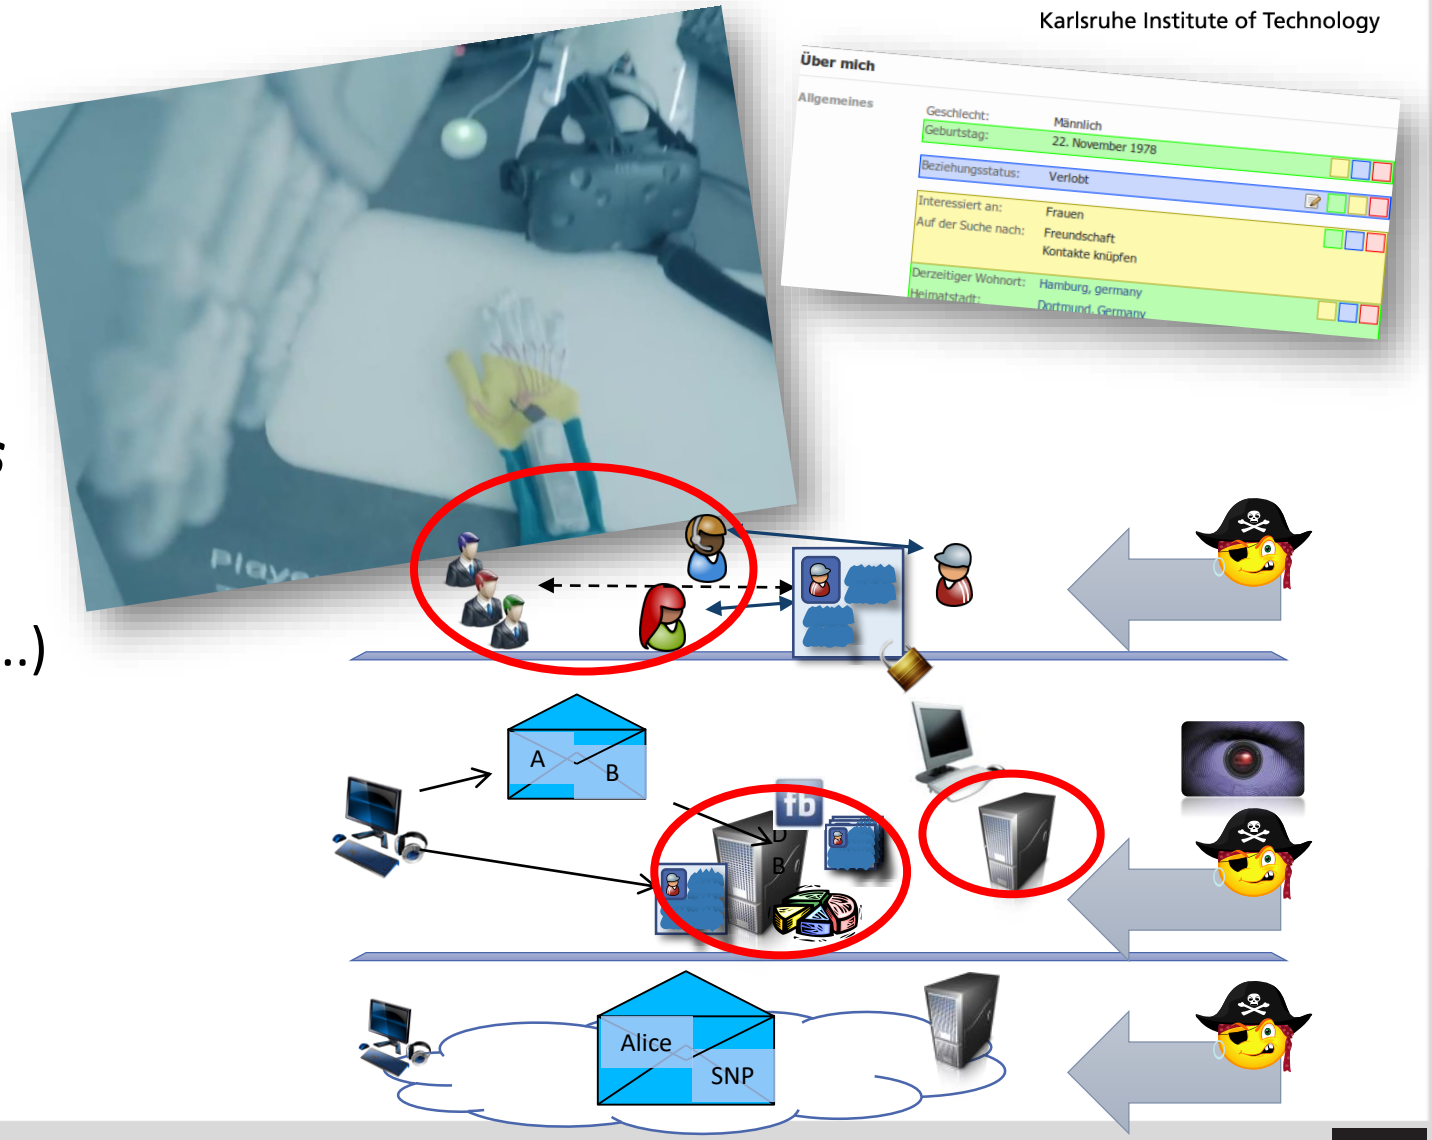

What we're working on…

- ***User understanding***
  - Intention recognition
  - Privacy analyses
  - Data sanitation

- ***Privacy-Enhancing Technologies***
  - Anonymity metrics
  - Anonymous services (f2f/Web)
  - Anonymous Communication (Tor, ..)

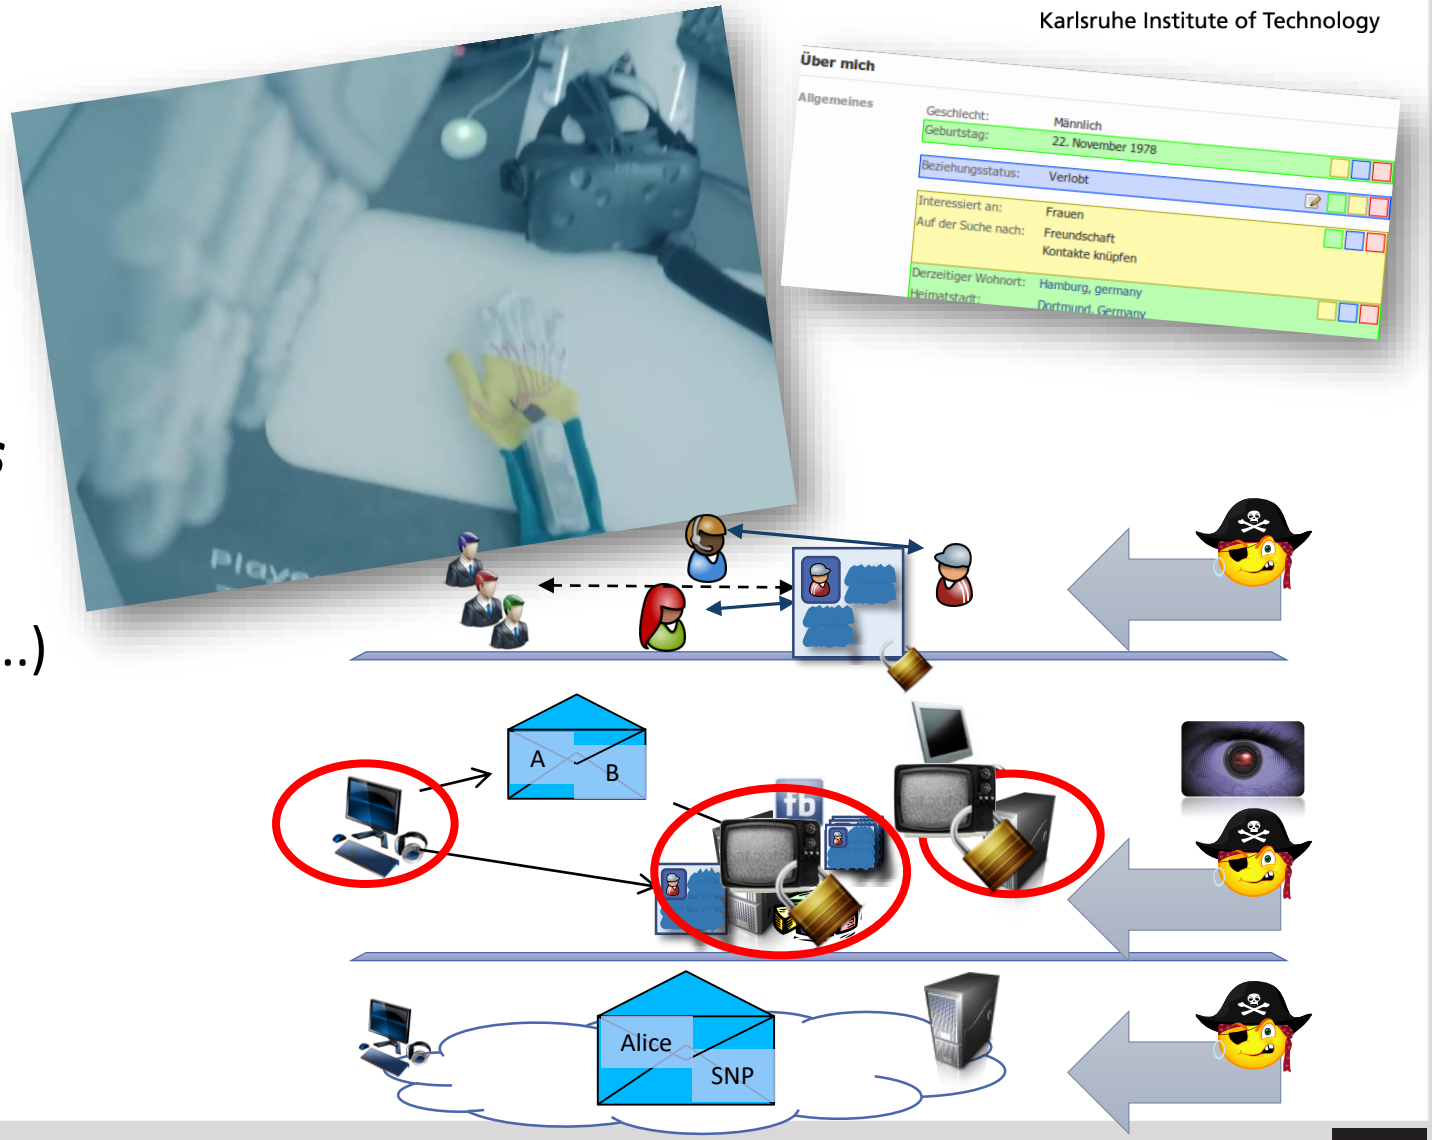

Anonymity Notions

- Plethora of anonymizers around
  - *TOR, AN.ON, DC, HORNET, Loopix, ZCash,…*
  - Claim *„Sender-Anonymity“*, or *„Recipient-Anonymity*, or *„Transaction Confidentiality“*
  - Literature defines *„Unlinkability“*, *„Unobservability“*, *„Pseudonymity“*, *„*-Anonymity“*, *„Anonymity Sets“*, *„Indistinguishability“*
  - So what does all this actually mean?

- Study
  - Game-based formalization of anonymity online
  - Consider all communication properties
  - Define and analyse privacy notions and their dependencies, rigorous protocol analysis

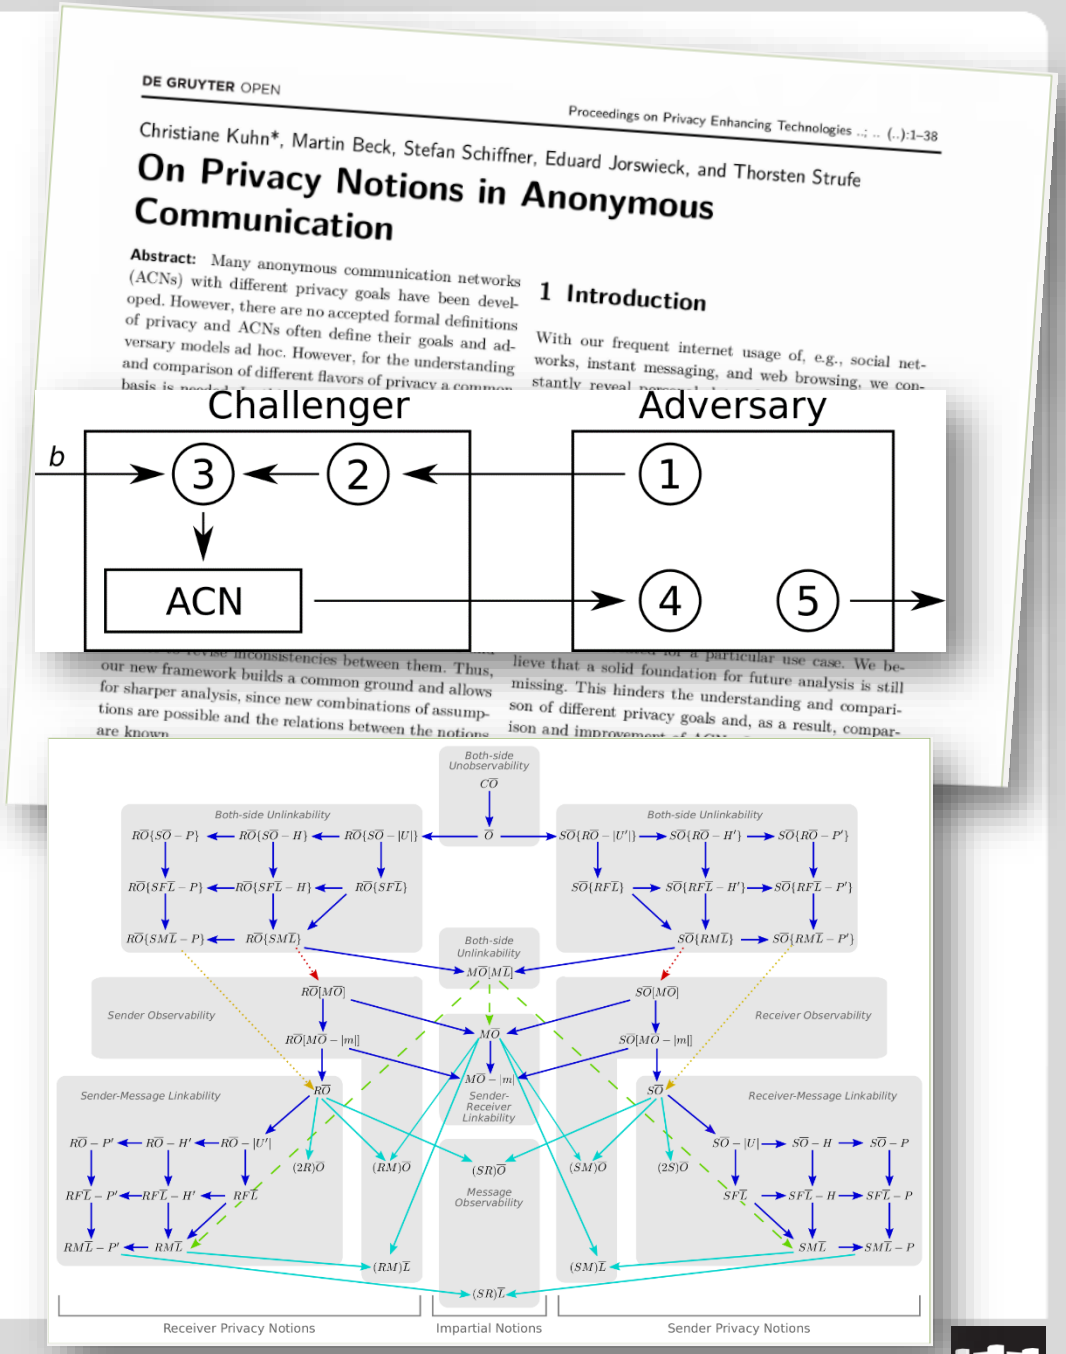

[17] PETS ’19
[18] S&P ‘20

What we're working on

- ***User understanding***
  - Intention recognition
  - Privacy analyses
  - Data sanitation

- ***Privacy-Enhancing Technologies***
  - Anonymity metrics
  - Anonymous services (f2f/Web)
  - Anonymous Communication (Tor, ..)

- ***Network security***
  - Network isolation, VPNs
  - 5G/6G security (now also: architectures)
  - PHYsec (now also with quantum ;-))

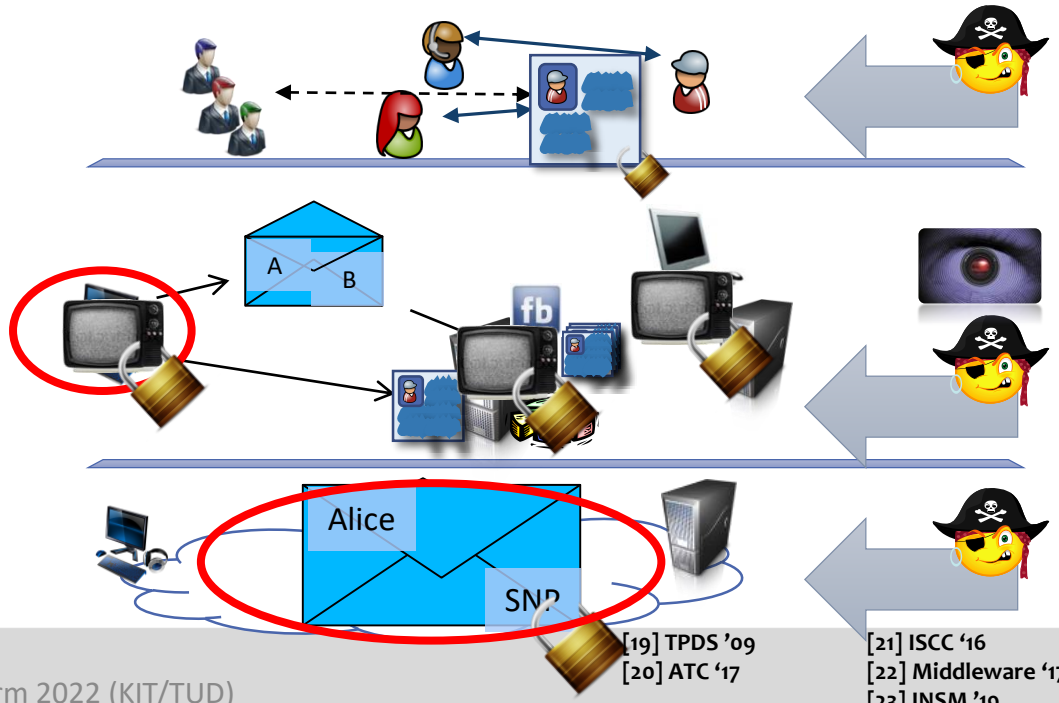

[19] TPDS '09
[20] ATC '17
[21] ISCC '16
[22] Middleware '17
[23] INSM '19

Karlsruhe Institute of Technology

KASTEL

Resilient Networking

Lecture/Reading group Winter term 2022

Some Words Regarding this Course

- Main topic of the course is the *security of deployed, crucial networks, networking functions, and network protocols.*

- Considering the Internet: *networking is an essential service, hence the networking infrastructure is/may be the main target of attacks!*

- ***Now what!?***

Preliminary Course Overview

1. Introduction
2. Graphs and graph theory
3. Crypto basics (Symmetric/Asymmetric/MACs)
4. *Link-Layer Security*
5. Resilient Routing (Attacks on BGP, SBGP)
6. *IPsec*
7. *TLS*
8. DNS Security
9. DDoS and Countermeasures
10. *Resilient Overlay Networks / Blockchain / Darknets*
11. *Intrusion Detection and Response*

Organizational Matters

- There will be some ex-cathedra parts, but please ask and discuss as much as possible!

- Course Language
  - Slides are in English, presentation as you prefer
  - => What‘s your language of preference?

- Slide history
  - Based on several former courses given at TU Ilmenau, Uni Mannheim, TU Darmstadt, and Dresden
  - Heavily derived from „Network Security“ and „Protection of Communication Infrastructures“ of/with Prof. Schäfer in Ilmenau and extended with Prof. Fischer‘s input from UHH

Material

- Slides will be on the Web site

- Literature/References
- Schäfer, Roßberg: Network Security
- For crypto: Dan Boneh‘s coursera course

- David Kahn: The Codebreakers
- Simon Singh: The Code Book

Organizational matters

- Lecture
- Fri 9:45 – 11:15
- 50.34 : 301

- Exercises
- Tue/Thu 14:00 – 15:30
- 50.34 252 (first meeting in CW 47: Nov 22, start preparing)

- Exams
- Oral exams, make appointments
- Procedure:
  - Questions available in German (and English upon request)
  - Answers given in German (and English upon request)

- All necessary information (will be) on the Web site: https://ps.tm.kit.edu/english/139_600.php *(ILIAS/OPAL?)*
- *Register to the mailing list* ***dud-resnets@groups.tu-dresden.de*** *!*

The Reading Group (Exercises)

- Exercise course will be organized as a reading group
  - Papers (links) available on the webpage (soon)
  - Read papers early…
  - One paper with relation to lecture topics will be presented (by a random ***one*** of ***you***!) and discussed (by ***you***!) each week (please take note of the emphasize on ***YOU :-)***)

The Reading Group

Intention of the reading group is to learn

- from good (and bad) scientific papers
- how to stay up to date and inform yourselves at the source
- that what others do is mostly no rocket science
- how to read a paper properly
  (probably not in the order from beginning to the end!)

Different kinds of papers

- Papers: the classic form of scientific content spreading, a single contribution
  - *Workshops*: Early ideas, WiP, Challenges/discussions ("*Recurring issues with spark-plug electrodes*")
  - *Conferences*: concise studies ("*On the electrode shapes in spark-plug design*")
- Journal articles: self-contained ("*On spark-plug design*")
- Surveys: summarizing a field or research area

The Reading Group – Reviewing Papers

- Paper idea
  - What is the field of research?
  - What is the motivation of the paper?
  - What is the problem the paper tries to solve?
  - What is the exact research question?
  - What is (are) the paper hypothes(i|e)s?
  - How relevant is this research?

- Paper content
  - What is the ***claim***, what are the ***assumptions*** of the paper?
  - Which definitions are contained?
  - What is the idea for solving the problem?
  - Which implications does it entail?
  - How is the evaluation carried out? Does it suffice to demonstrate/substantiate the claims? What about the results?

- Critical acclaim: Merits & Shortcomings

With a little help by a random stranger…

Paper	Title, Author(s)
Field of research	
Exact research question	
Relevance (Claim)	
Hypothesis	
Content	Assumptions
	Definitions
	Overview of solutions
	Evaluation style, procedure, results
Crittical acclaim	Merits
	Shortcomings

Survey	Title, Author(s)
Field of Research	
Exact problem domain	
Content	Assumptions / Definitions
	Aspects, requirements, concepts, properties
	Classification
Critical acclaim	Sensibility of classes
	Completeness
	Merits
	Shortcomings

Questions?

Developing our terms…

What are „Resilient Networks“?

- *“Resilience is the ability of an object to spring back into shape”*
- *“Resilience is the ability of the network to provide and maintain an adequate level of service in the face of challenges to normal operation”*
- *“Resilience is the ability of the network to provide and maintain an acceptable level of security service in case some nodes are compromised.”*
- Challenges? Compromised nodes…?
- What kind of *problems, challenges, threats* could you imagine?
- What ***exactly*** do these terms mean, anyway?

Sterbenz: Resilinets, 2006-2008
Chen et al.: Sensor Network Security, 2009

Resilience Disciplines

Resilience comprises a multitude of disciplines

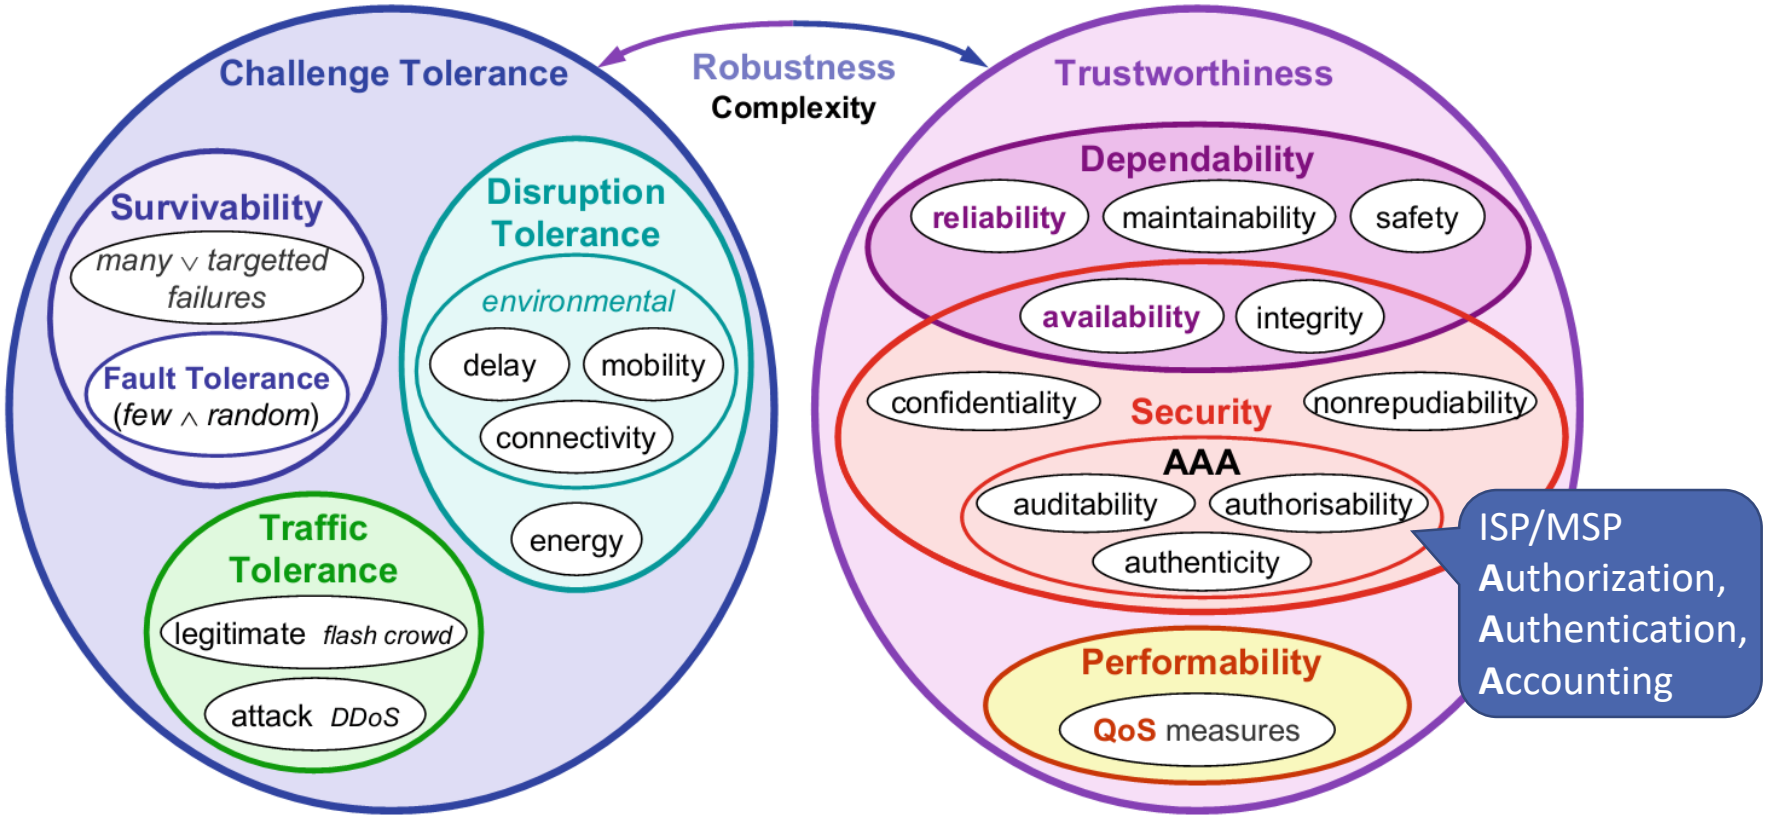

Sterbenz, James P.G., Hutchison, David, Çetinkaya, Egemen K Jabbar, Abdul, Rohrer, Justin P, Schöller, Marcus and Smith, Paul. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. IEEE Computer Networks, 2010

Resilience – Challenge Tolerance (1)

- ***Survivability***
  - Capability of a system to fulfill its mission,
    - in a timely manner,
    - in the presence of ***threats*** such as ***attacks*** or large-scale natural ***disasters.***
  - Covers ***correlated failures*** as result of ***intelligent adversary*** and failures of large parts of network infrastructure
  - Requires **diversity**: same fate unlikely to be shared by parts of system undergoing correlated failures

- ***Fault tolerance***
  - Subset of survivability
  - Ability of system to tolerate faults to ***prevent service failures***
  - Relies on **redundancy** to compensate random uncorrelated failures of components
  - Provides no sufficient coverage when facing correlated failures

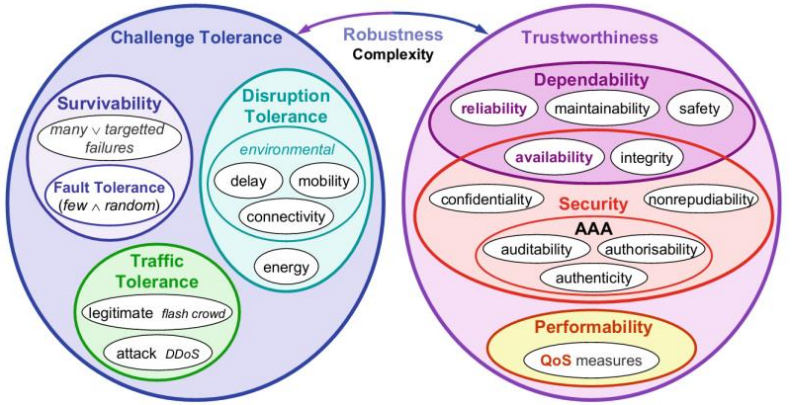

Resilience – Trustworthiness (1)

- ***Dependability***
  - Quantifies resilience of the service delivery by a system
  - Basic measures
    - Mean Time To Failure (MTTF)
    - Mean Time To Repair (MTTR)
  - Consists of
    - ***Availability***: readiness for usage
    - ***Reliability***: continuous service delivery

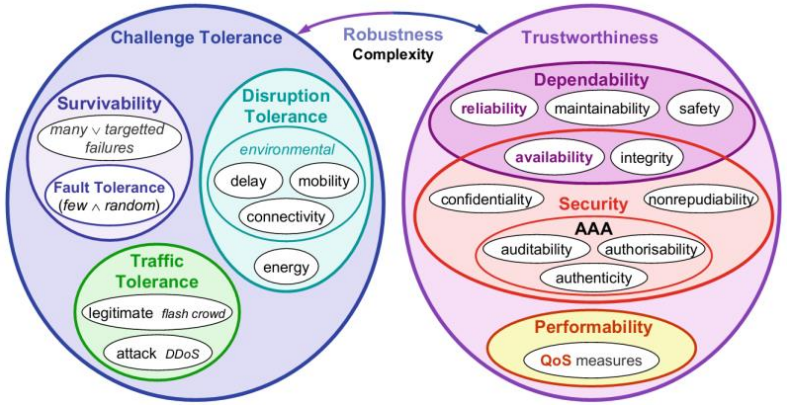

- ***Security***
  - Property of a system, and the measures taken such that it protects itself from ***unauthorized access or change***
  - Security shares availability with dependability
  - However, we assume a strategic adversary (worst case, repeatedly…)

Karlsruhe Institute of Technology

KASTEL

The Security in Resilience…

Introducing Actors of the Play

- For clarity it‘s good to have some model…
- The classic security – scenario:

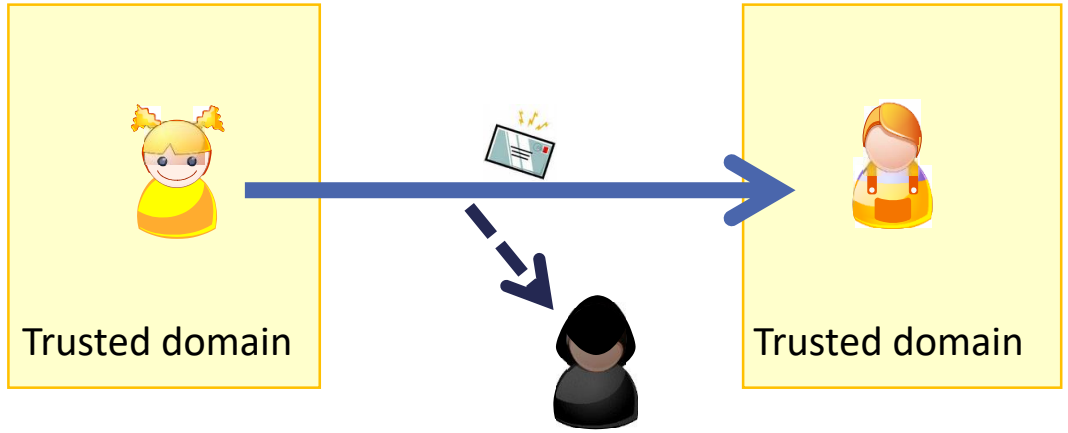

- The RN – scenario:

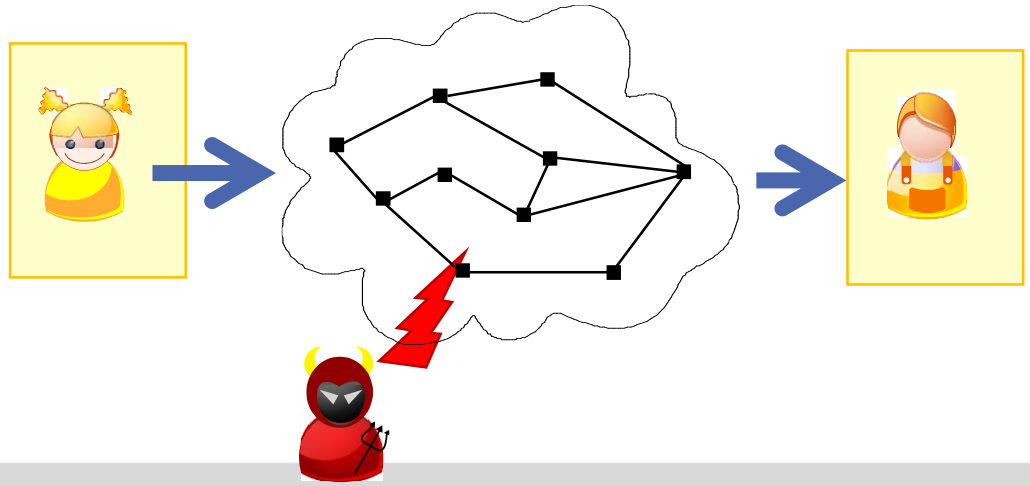

Threats in Communication Networks

- Abstract Definition:
  - A threat is any possible event or sequence of actions that might lead to a violation of one or more security goals
  - The actual realization of a threat is called an ***attack***

- Examples:
  - A hacker breaking into a corporate computer
  - Disclosure of emails in transit
  - A hacker temporarily shutting down a website
  - Someone using services or ordering goods

    in the name of others
  - ...

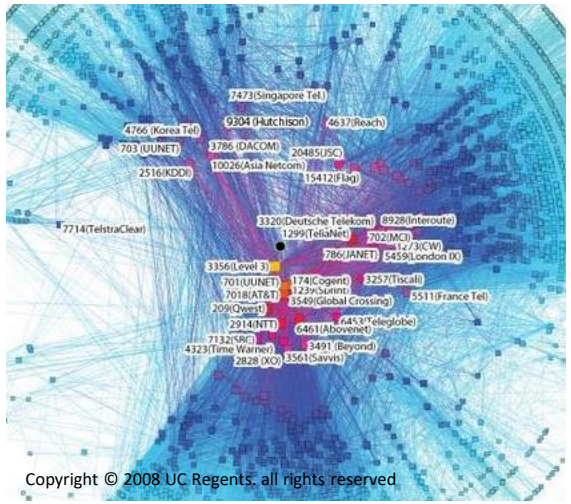

Potential Attackers and an Adversary Model

A word on assumptions.

- Assume an omnipotent adversary. She could:
- access all information of interest
- compromise arbitrary intermediate systems
- physically destroy any or all components

- Could we deal with this?
- Unfortunately, no:
- *„Nothing can protect from an omnipotent adversary.“*

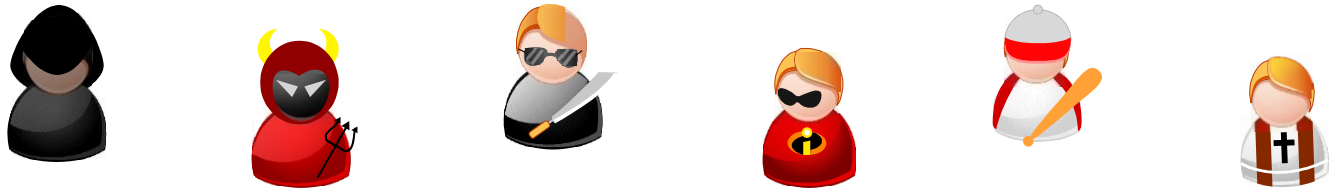

- More realistic (specific!) model of adversaries needed.

On Eve, Mallory, Craig, and Trudy…

- An ***adversary model*** needs to define
- The *intention* of the adversary
  - Break and/or access <something>
- The *behavior*
  - Passive or active?
- The *capabilities* of an attacker
  - Computational capacity (often: think complexity class)
  - Resources (time and money)
- The *area of control*
  - Insider or outsider?
  - Local, regional, or global?

A little exercise for the weekend: what are the adversary models in specific examples:

https/TLS? Email-Encryption? TOR?

The Dolev - Yao Model

- Mallory has full control over the communication channel
  - Intercept/eavesdrop on messages (passive)
  - Relay messages
  - Suppress message delivery
  - Replay messages
  - Manipulate messages
  - Exchange messages
  - Forge messages

- But:
  - Mallory ***can't*** break (secure) cryptographic primitives!

Threats Technically Defined

- *Masquerade:*
  - An entity claims to be another entity
- *Eavesdropping:*
  - An entity reads information it is not intended to read
- *Authorization violation:*
  - An entity uses a service or resources it is not intended to use
- *Loss or Modification of (transmitted) information:*
  - Data is being altered or destroyed
- *Denial of Communication Acts (Repudiation):*
  - An entity falsely denies its participation in a communication act
- *Forgery of information:*
  - An entity creates new information in the name of another entity
- *Sabotage:*
  - Any action that aims to reduce the availability and / or correct functioning of services or systems

Security Goals in Application Environments

- Public Telecommunication Providers:
  - Protect subscribers' privacy
  - Restrict access to administrative functions to authorized personnel
  - Protect against service interruptions
- Corporate / Private Networks:
  - Protect corporate confidentiality / individual privacy
  - Ensure message authenticity
  - Protect against service interruptions
- All Networks:
  - Prevent outside penetrations (who wants hackers?)
- Security goals are also called ***security objectives***

Security Goals Technically Defined (CIA)

- ***C**onfidentiality:*
  - Data transmitted or stored should only be revealed to the intended audience
  - ***Confidentiality of entities*** is also referred to as ***anonymity***
- *(Data) **I**ntegrity:*
  - It should be possible to detect any modification of data
  - This requires to be able to *identify* the creator of some data
- ***A**vailability:*
  - Services should be available and function correctly

- *Accountability:*
  - *It should be possible to identify the entity responsible for any communication event*
- *Controlled Access:*
  - *Only authorized entities should be able to access certain services or information*

Several other models have been proposed, anything beyond CIA is constantly subject to arguments and discussions…

Interlude: Security Services

- *Security Service:*
  - An abstract “service” seeking to ensure a specific security property
  - Can be realised with the help of ***cryptographic*** algorithms and protocols or with ***conventional*** means:
    - Keep electronic document on a floppy disk confidential by storing it on the disk in an encrypted format ***or*** locking away the disk in a safe
    - Usually a ***combination*** of cryptographic and other means is most effective

Security Services – Overview

- *Authentication*
  - Ensure that an entity has in fact the identity it claims to have
- *Integrity*
  - Ensure that data created by specific entity isn't modified ***without detection***
- *Confidentiality*
  - Ensure the secrecy of protected data
- *Access Control*
  - Ensure that each entity accesses only services and information it is entitled to
- *Non Repudiation*
  - Prevent entities participating in a communication exchange from later falsely denying that the exchange occurred

That was fairly abstract…

How can we operationalize this?

Network Security Analysis

- To find countermeasures, threats have to be evaluated appropriately for a given network configuration.

- Therefore, a detailed network security analysis is needed that:
  - evaluates the ***risk potential*** of the general threats to the entities using a network, and
  - estimates the ***expenditure*** (resources, time, etc.) needed to perform known attacks.

→ Attention: *It is generally impossible to assess unknown attacks!*

Architectural View of the Threatened “Object”

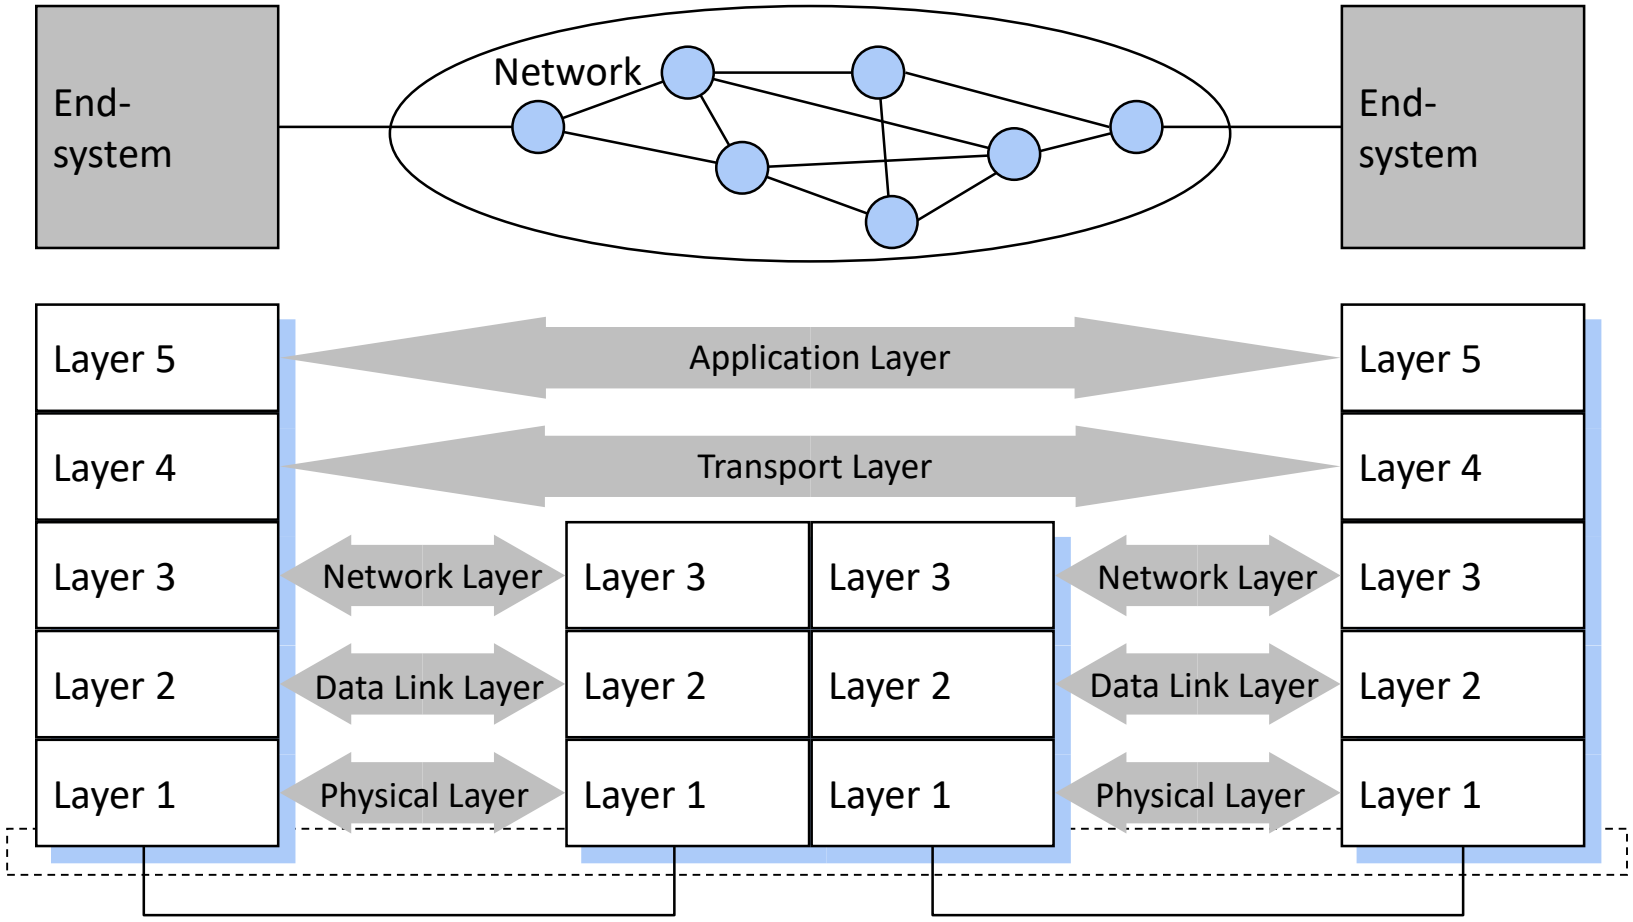

Communication in Layered Protocol Architectures

Security Analysis of Layered Protocol Architectures 1

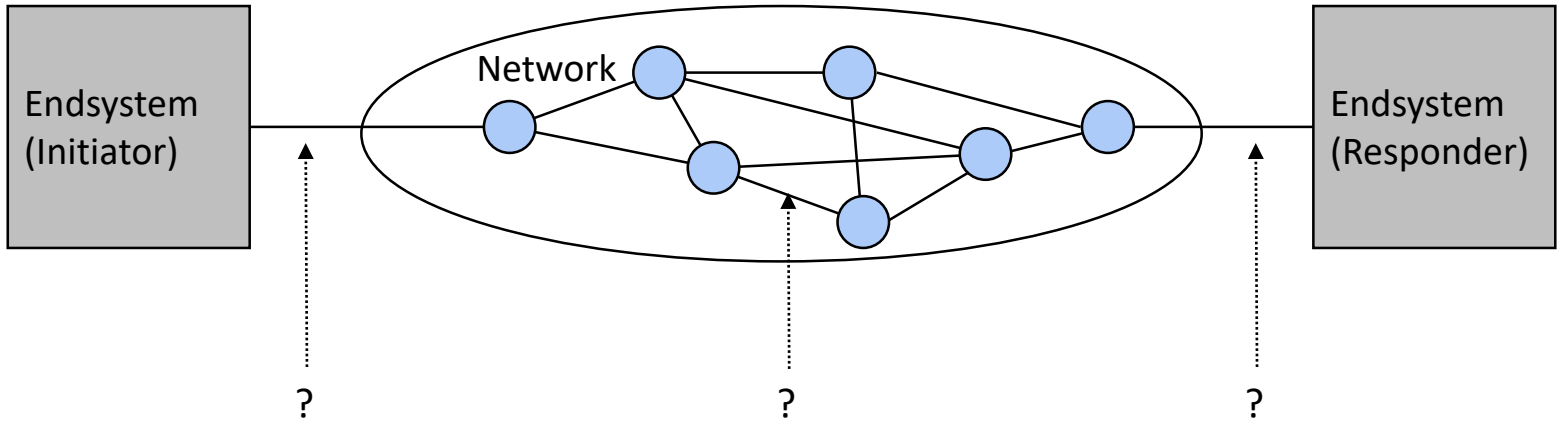

Dimension 1: At which interface could an attack take place?

Security Analysis of Layered Protocol Architectures 2

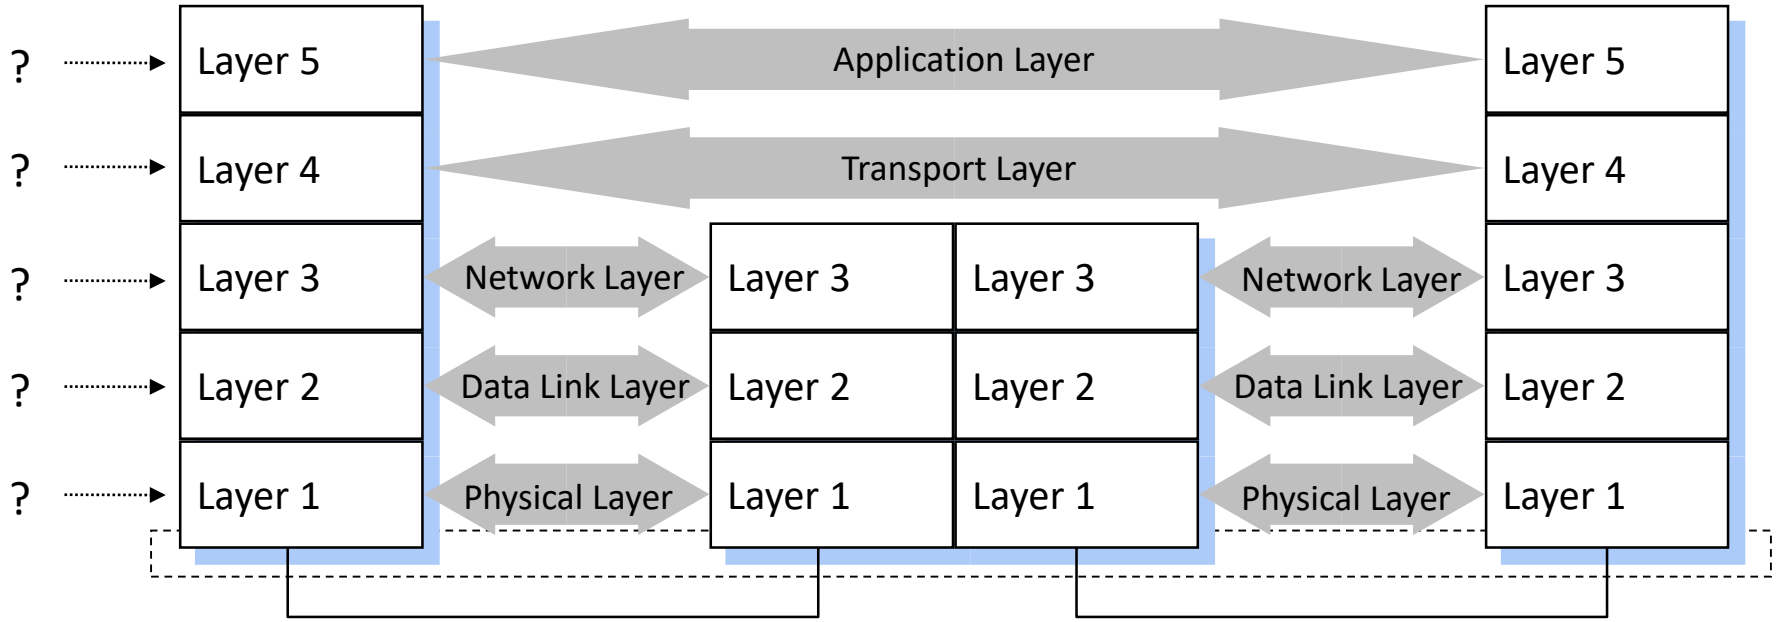

Dimension 2: In which layer could an attack take place?

Potential Points of Attack

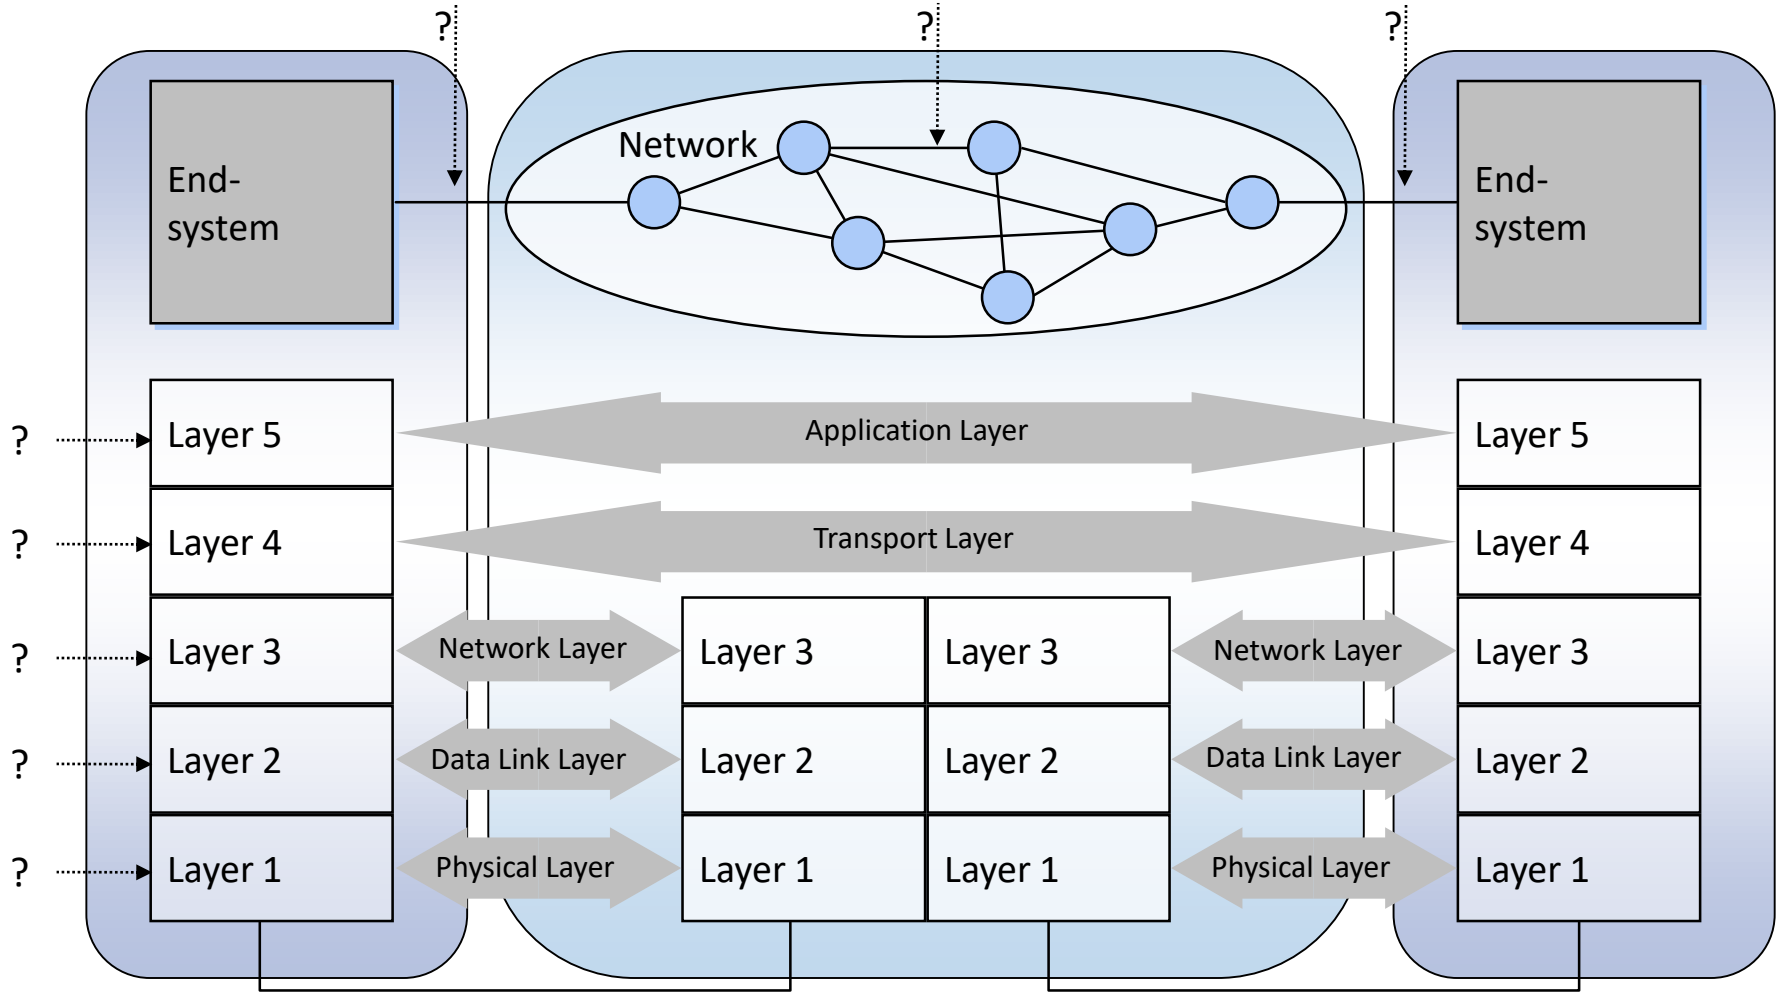

Towards a Systematic Threat Analysis

- One approach: produce arbitrary threat list by any ad-hoc brainstorming method

- Example: Hospital Information System
  - Corruption of patient medical information
  - Corruption of billing information
  - Disclosure of confidential patient information
  - Compromise of internal schedules
  - Unavailability of confidential patient information
  - …

Downside: not very systematic…

- Drawbacks of this approach:
  - Questionable completeness of identified threats
  - Lack of rationale for identified threats other than experience
  - Potential inconsistencies (e.g. disclosure vs. unavailability of confidential patient information in the example above)

Approaches for Systematic Threat Modeling

- Explicit quantification of security is hard (impossible?)
- Threat modelling is a soft task

- Alternative management approaches have been suggested
  - STRIDE
    - Risk identification (Microsoft: Kohnfelder and Garg, 1999)
    - ***S**poofing, **T**ampering, **R**epudiation, **I**nformation-disclosure, **D**oS, **E**levation of Privilege*
  - DREAD
    - Risk assessment, as used e.g. by OpenStack (among others)
    - ***D**amage, **R**eproducibility, **E**xploitability, **A**ffected Users, **D**iscoverability*
  - Threat Trees (Amoroso, 1994) (later on: „Attack trees“)

Shostak: “Experiences Threat Modeling at Microsoft”, 2004

Threat Trees: One Systematic Threat Analysis Approach

Karlsruhe Institute of Technology

- A *threat tree* is a tree with:
  - *nodes* describing threats at different levels of abstractions, and
  - *subtrees* refining the threat of the node they are rooted at,
  - where the child nodes of one node give a *complete refinement* of the threat represented by the parent node

- Technique for establishing threat trees:
  - ***Start*** with general, abstract description of complete set of threats for a given system (e.g. “security of system X compromised”)
  - ***Iteratively***, gradually introduce detail by carefully refining the description
  - Each node becomes root of a ***subtree*** describing ***threats*** represented by it
  - Eventually, each ***leaf node*** of the tree provides a description of a threat that can be used for a (less arbitrary) threat list

- The main idea of this technique is to postpone the creation of (arbitrary) threat lists as much as possible

Example: A Hospital Information System Threat Tree

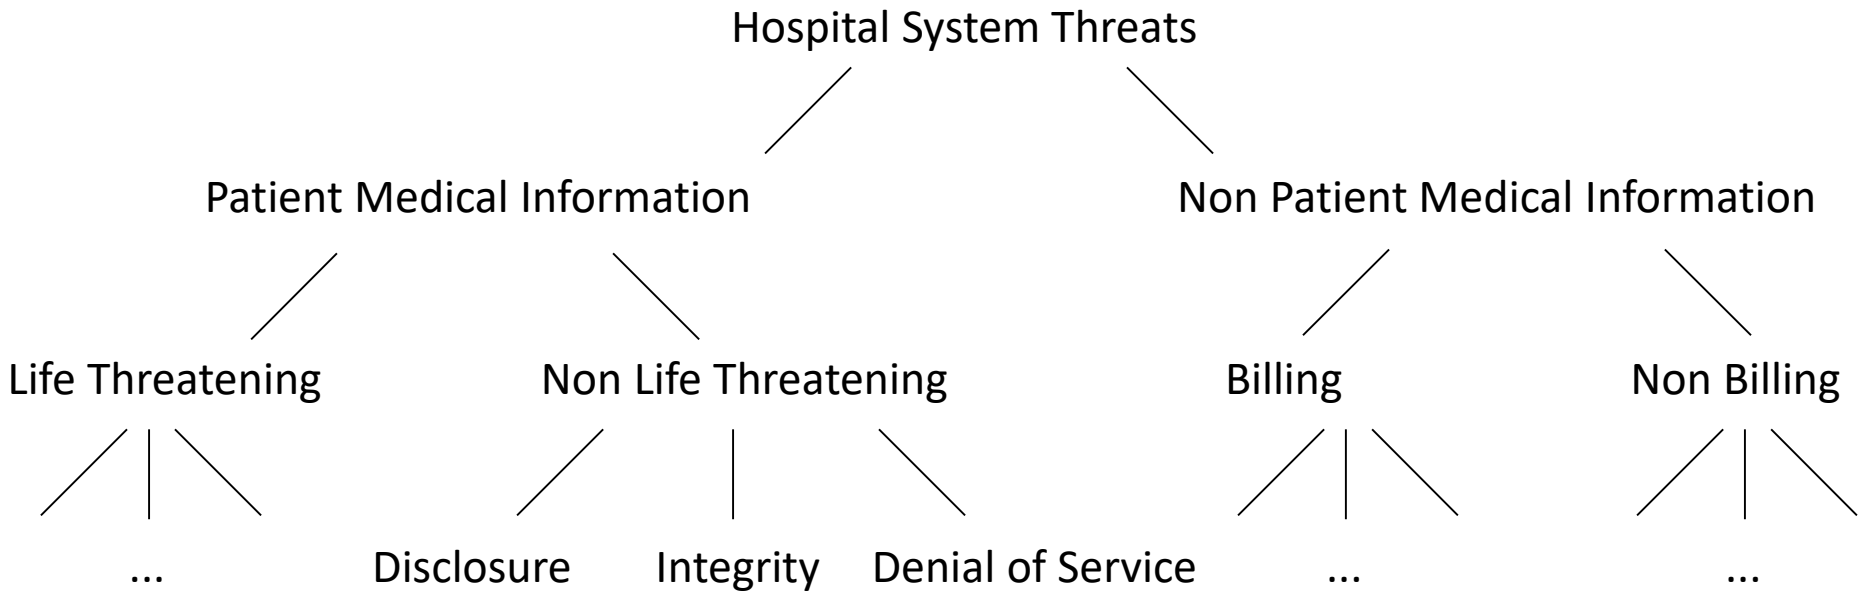

➡ At each level of refinement the child nodes of a node ***must*** maintain *demonstrable completeness* to allow for confidence that nothing is missing

(source: [Amo94])

Inferring Composed Threat in Threat Trees

- Child nodes can have different relations to their parent nodes
- The two most common relations are AND and OR:

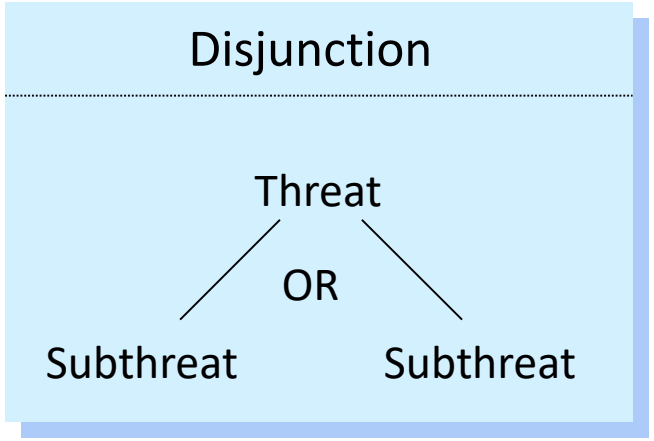

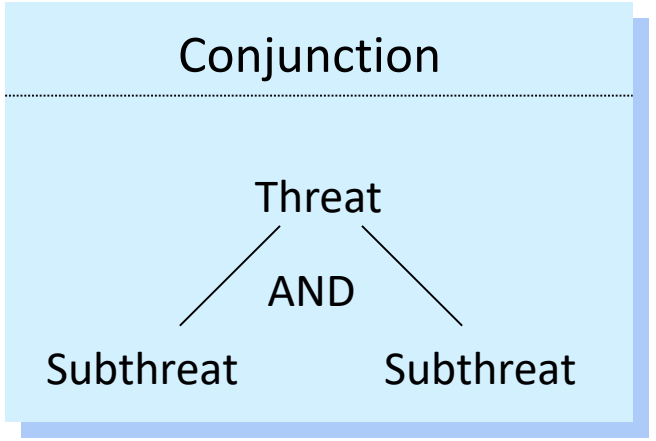

- These relations can be used to infer composed threat:
  - Augment nodes with effort estimations (e.g. easy, moderate, high)
  - OR-related composed threat inferred as the lowest effort value subtree (the attacker will most likely take the easy way...)
  - For conjunction, infer highest effort (all threats have to be realized)

Risk-Assessment/Quantification with Threat Trees

- Appropriate attributes are, e.g., estimated criticality and attacker effort for individual threats
- Threat trees then can help to gain insight where to spend resources to decrease the overall system's vulnerability:

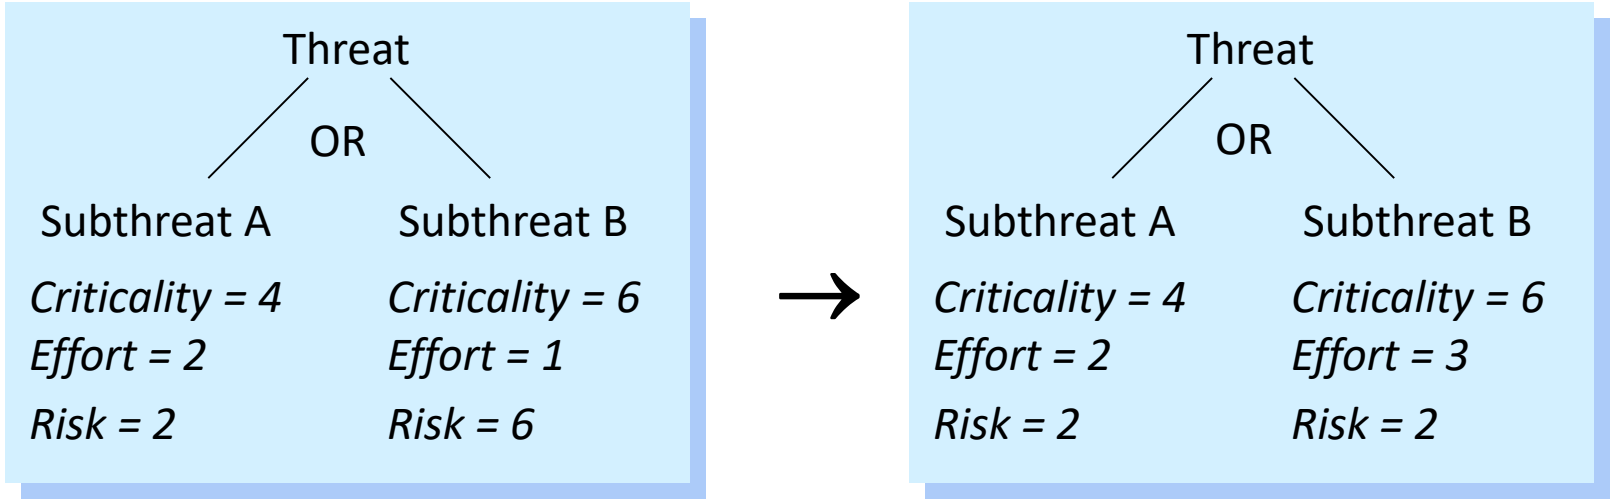

- The second threat tree re-evaluates the risk after some protective measure has been taken to increase the attacker's effort for subthreat B
- Here, risk is assessed as: ***Risk = Criticality / Effort***

Variation of the Game: Attack Trees

- NSA/Darpa/Schneier‘s approach:
  - Model the attacker‘s goal as root node
  - Branches model means of reaching the goal
  - Leaf nodes enumerate specific attacks

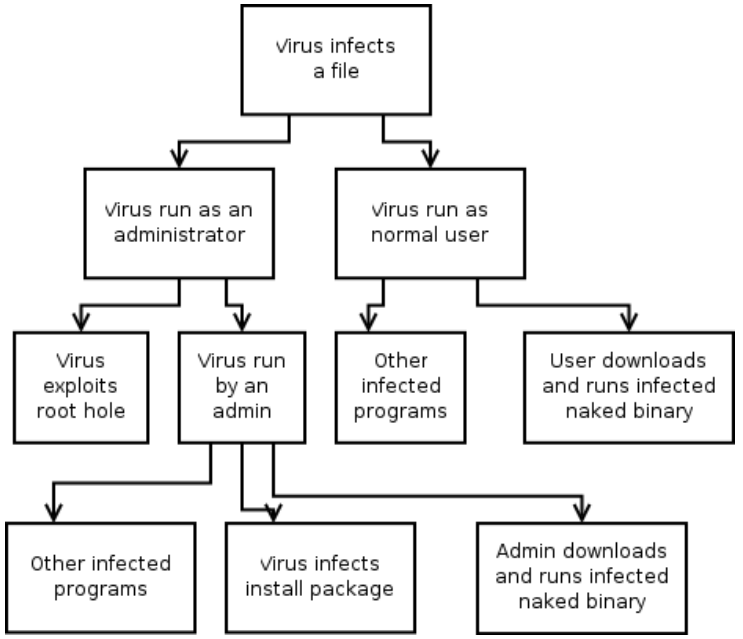

Source: wikipedia

Summary (High Level System Security Engineering Process)

- Specify system architecture:
  - Identify components and interrelations
- ***Identify threats***, vulnerabilities and attack techniques:
  - The threat tree technique provides help for this step
- ***Estimate*** component ***risks*** by adding attributes to the threat tree:
  - However, removing subjectivity from initial assessments is often impossible and other attributes than criticality and effort (e.g. risk of detection) might have to be considered as well
- Prioritize ***vulnerabilities***:
  - Taking into account the components' importance
- Identify and install ***safeguards***:
  - Apply protection techniques to counter high priority vulnerabilities
- Perform potential ***iterations*** of this process
  - Re-assess risks of the modified system and decide, if more iterations are required

Countering Attacks: Three Action Classes

- ***Prevention**:*
  - Measures taken to avert that an attacker succeeds in realizing a threat
  - Examples:
    - *Cryptography*: encryption, computation of modification detection codes, running authentication protocols, etc.
    - *Firewalls*: packet filtering, service proxying, etc.
- ***Detection**:*
  - Measures taken to recognize an attack *while or after it occurred*
  - Examples:
    - Recording and analysis of audit trails
    - On-the-fly traffic monitoring
- ***Reaction**:*
  - Measures taken in order react to *ongoing (mitigation and healing) or past attacks*
  - *Examples:*
    - *Adding new firewall rules*
    - *Traffic re-routing*
- *(DDS: Prevention, Removal, Forecasting, Tolerance/Graceful degradation)*

Course Objectives

- This course tackles the following aspects:
  - Threats to and measures for ensuring ***availability***
  - Threats and measures concerning systems (beyond pure network security protocols which are more targeting transmission security)
  - Measures for intrusion detection and response

- *Considering the Internet: networking is an essential service, hence the networking infrastructure is/may be the main target of attacks! We'll hence be looking at the security of deployed, crucial networks, networking functions, and network protocols.*

Summary

- You know who we are
- You know what to expect from the lecture
- You have seen some trends that are happening
- You have been introduced to Alice, Bob, Eve, and Mallory
- You understand what threats are … and what this means
- You can tell security goals (*CIA!*) from security services
- You know how to perform a network security analysis using threat trees ;-)

Questions?