


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

Prof. Dr.-Ing. Georg Carle and 18 research staff
Organization: Daniel Raumer
Contact: seminar@net.in.tum.de

Challenging Topics!
Sometimes previous knowledge
required!



Uhrenturm der TUM

Content



www.fotoila.de

Administrative Things for all Seminars

- Basic Information
- Grading
- Responsibilities

Topics

Basic Information

- Lecturer:
 - Prof. Dr.-Ing. Georg Carle
- Organization: seminar@net.in.tum.de (only use this mail address!)
 - Daniel Raumer
- Overview
 - Main Language: German
 - we will offer an English track (presuming a minimum of 4 participants)
 - Extent: 2 SWS, (5 ECTS)
 - 5 ECTS * 30 hours = 150 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...)
- This means: if there is an English only track, there are two tracks instead of one. We publish two different schedules. You just have to attend talks listed on your schedule.

Seminar Procedure

- First version of your paper
 - Agree on the content with your advisor
 - Use the supplied paper template from the Moodle course
 - Keep in touch with your advisor
 - Try to finish well in time so you advisor can give you feedback
- Present your topic
- 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
- Paper extent
 - Your paper **MUST** be 6-8 pages in ACM 2-column style (including references etc.)

FI vs IITM

Both seminars offer similar topics

- Again, check the proceedings (link on previous slide)

Future Internet

- Topic presentation: Tuesday, August 1st, 14:00, room 03.07.023
 - This is just one week 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 (October 12/13th)
- Registration limited to 16 participants

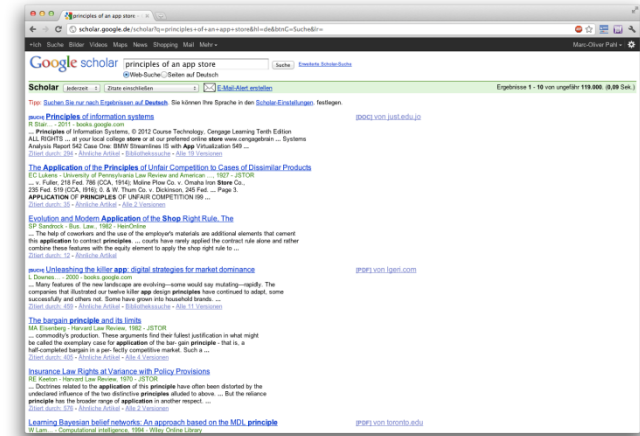
Innovative Internet Technologies and Mobile Communications

- Topic presentation: Monday, October 16th, 14:00, room 03.07.023
- Seminar paper to be written during the semester (until summer holidays)
- Weekly sessions on Monday, 14:00, starting in December of January 2017/18 (probably 4 meetings)
- Registration limited to 14 participants

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
 - ...



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.

Your task is to do research and write a scientific text about a specific topic beyond basic lecture content.

→ Your advisor is not responsible for your 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

- Prepare your talk
 - Finished slides must be discussed with advisor 1 week before the talk
 - Advisors usually offer the opportunity of rehearsal talks
- Present your work
 - Scientific talk; build on knowledge from GRNVS and other basic lectures
 - Do not present your paper; present the main results & give an interesting talk
- Session chair for one talk
 - Introduce the speaker
 - 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 in advance
 - Attending the talks is mandatory for passing the course

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%)

Grading – Influencing Factors (1)

- First version of paper must be acceptable
 - Grade worse than 4.0 → disqualification (seminar graded as 5.0)
 - Less than 6 pages in the seminar template → disqualification
(Not 5,1 + empty space; Not 8 + references)
- 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

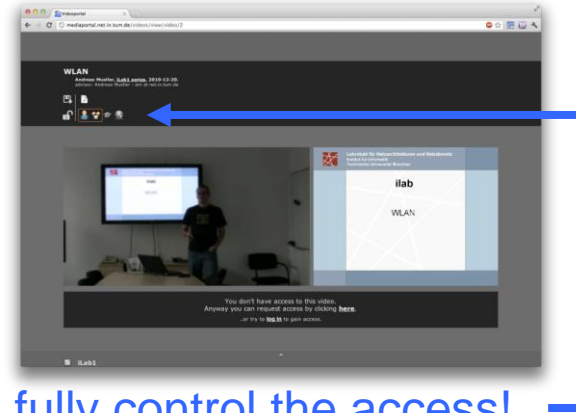
Grading – Influencing Factors (2)

- 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 will be reported to the examination office
 - Summary when and why to cite:
<http://oxford.library.emory.edu/research-learning/citation-plagiarism/citing.html>
- Unauthorized absence during talks
 - 0.3 degrading per missed talk on your presentation grade
 - Talk graded worse than 4.0 → disqualification (seminar graded as 5.0)

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.



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

Get feedback from your colleagues (not graded)

- Feedback forms to be filled out during the talks

Presentation Feedback Form

Title of presentation: _____

	Very poor / Confusing	Poor	Reasonable / OK	Good / Helpful	Excellent
Presentation skills	○○○○○	○○○○○	○○○○○	○○○○○	○○○○○
Style of presentation	○○○○○	○○○○○	○○○○○	○○○○○	○○○○○
Time keeping	○○○○○	○○○○○	○○○○○	○○○○○	○○○○○
Tone/ volume/speed	○○○○○	○○○○○	○○○○○	○○○○○	○○○○○
Use of presentation media	○○○○○	○○○○○	○○○○○	○○○○○	○○○○○
Interaction with audience	○○○○○	○○○○○	○○○○○	○○○○○	○○○○○
Content and structure	○○○○○	○○○○○	○○○○○	○○○○○	○○○○○

Seminar Proceedings

We give the opportunity to 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 register with your student card when leaving the room

- 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**
- We accept withdrawal only for special reasons
(e.g. not getting any of your topic preferences)

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



seminar@net.in.tum.de

Topics are part of research at our group

- Usually they can be continued as thesis, IDP, guided research, GEP, etc.
- 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
(more topics will be presented in the topic presentation session.)

Preview on topics that will be
offered in the next seminar

Resource management with learning algorithms

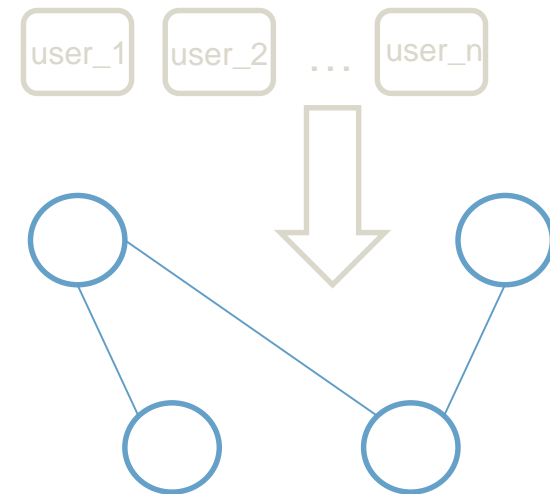
(Cora)

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?



Secure File Systems for the Cloud (Heiko)

You want to store your data on the Internet.

Secure File Systems

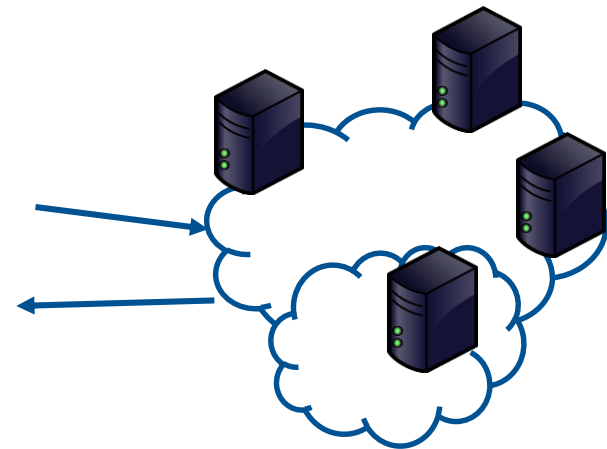
- Tahoe, ...

Secure File Systems using Network / Clouds

- DepSky
- Real ones...

Your Task:

- Idea of secure file systems
- How does transfer to networked file systems using Clouds?
- Give an example
- What's real?



[1] Wilcox-O'Hearn, Zooko, and Brian Warner. "Tahoe: the least-authority filesystem." *Proceedings of the 4th ACM international workshop on Storage security and survivability*. ACM, 2008.

{2} Bessani, Alysson, et al. "DepSky: dependable and secure storage in a cloud-of-clouds." *ACM Transactions on Storage (TOS)* 9.4 (2013): 12.

Comparison of Scientific Workflow Systems (Florian, Daniel)

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

Benefits:

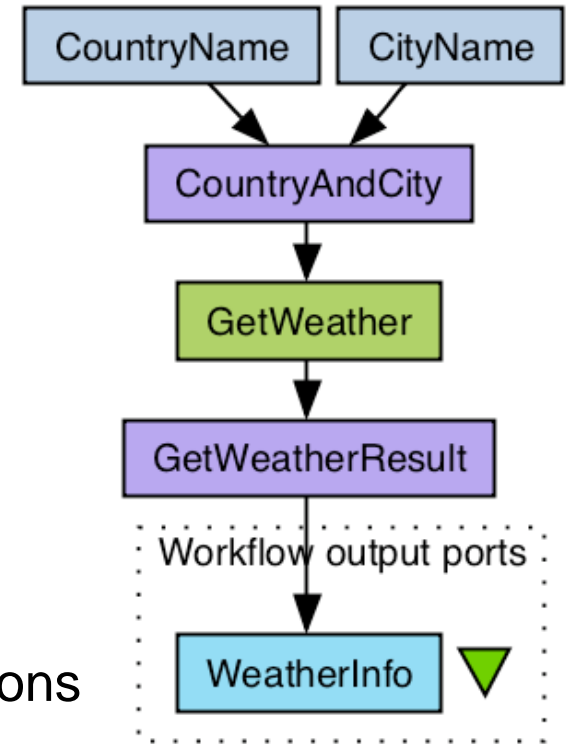
- Documented, reproducible process
- Simplified sharing/reuse between scientists
- Hide complexity

We operate a networking testbed

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

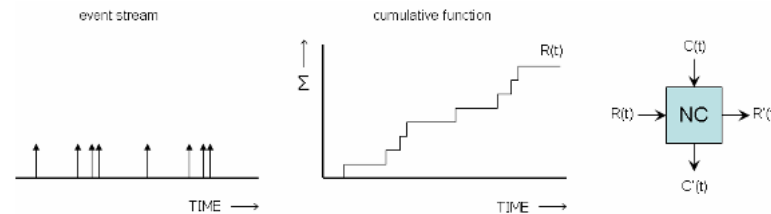
[1] V. Curcin et.al "Scientific workflow systems - can one size fit all?". 2008 Cairo Int. Biomedical Engineering Conference



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?*

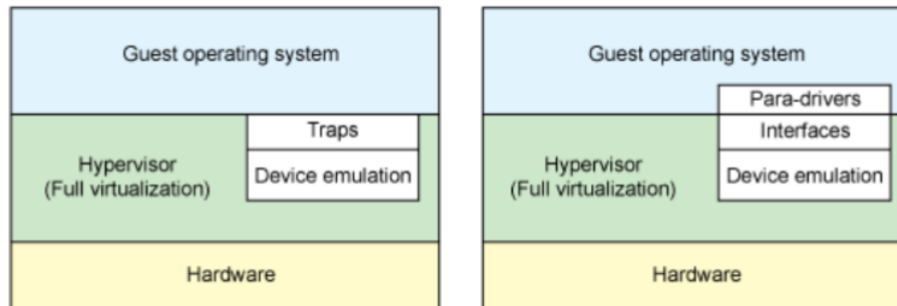
Starting point:

[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 Iguest", 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

Debian stretch has arrived... (Dominik, Sebastian)

... and with it comes Kernel 4.9. For a long time we have run experiments with kernel 3.16 or older, now it is time to upgrade.

It is not always easy to follow the newest developments and discussions regarding the Linux kernel, so this paper should give an overview and focus on the networking related changes from 3.16 to 4.9, which are worth for us to investigate their influence.

Your Task:

- What changed from 3.16 to 4.9?
- Focus on networking related topics
- Read changelogs
- Read official release messages, ...
- Optional practical work: create a Debian Stretch Live Image for our testbed



Stretch

(https://res.cloudinary.com/dli6s7kdg/image/upload/v1430288150/stretch-debian-lnx_svyegl.png)

Privacy Protection

in Network Anomaly Detection (Marcel, Stefan)

Anomaly Detection Systems are yet another system where a significant amount of data is piped through. While aiming for security and safety, data privacy is often not considered.

Nevertheless, privacy improvements or even full protection is possible in some use cases.

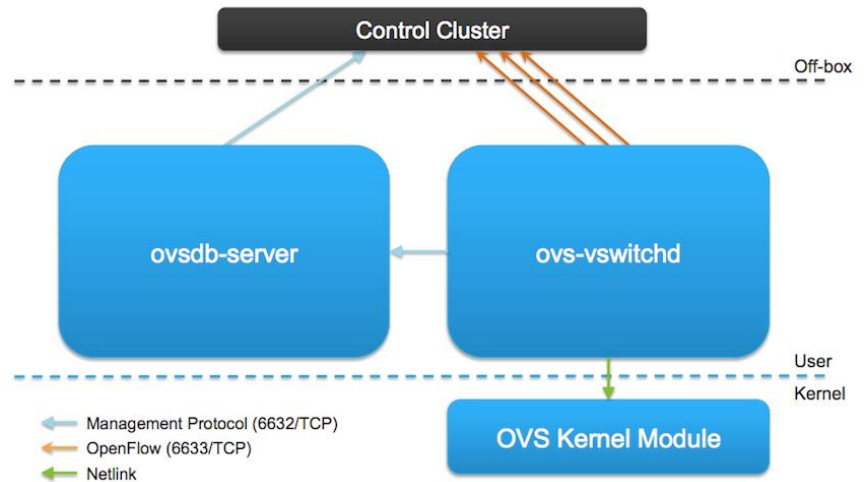
Your Tasks:

- Cooperate with the student of the previous topic
- Survey the state of the art
- Find out
 - Which protection strategies for anomaly detection exist?
 - How to protect network data (in arbitrary context)?
 - Does privacy-protection influence the performance of the detection system?
- Create an overview over possible methods and when to apply them

In-band SDN with Open vSwitches (OVS) (Sulovic)

Pure SDN architecture assumes decoupling of control and data planes, with all network intelligence centralized in SDN controller. The most commonly used protocol for control channel communication between the SDN controller and network infrastructure is OpenFlow.

Open vSwitch (OVS) is commonly used as an SDN switch. It can establish both, „out-of-band” and „in-band” control channel with the SDN controller.



Your Task:

- Describe the benefits of out-of-band and in-band control plane
- What are the prerequisites for in-band control plane establishment in SDN
 - How is the in-band control plane implemented in OVS
 - Which legacy protocols are used and what is their influence to the overall performance of SDN
- Describe how packets of legacy protocols (i.e., STP) are processed by OVS to bypass OpenFlow processing

Starting point:

[1] Pfaff, Ben, Justin Pettit, Teemu Koponen, Ethan J. Jackson, Andy Zhou, Jarno Rajahalme, Jesse Gross et al. "The Design and Implementation of Open vSwitch." In NSDI, pp. 117-130. 2015.

[2] www.openvswitch.org/

Federated Authentication and Access

Control Mechanism (Sree Harsha)

- Federation allows different autonomous systems to interoperate
- How to authenticate in federated environments?
- How to apply access controls in federated environments?
- Examples: Eduroam (WPS-EAP with Radius authentication)

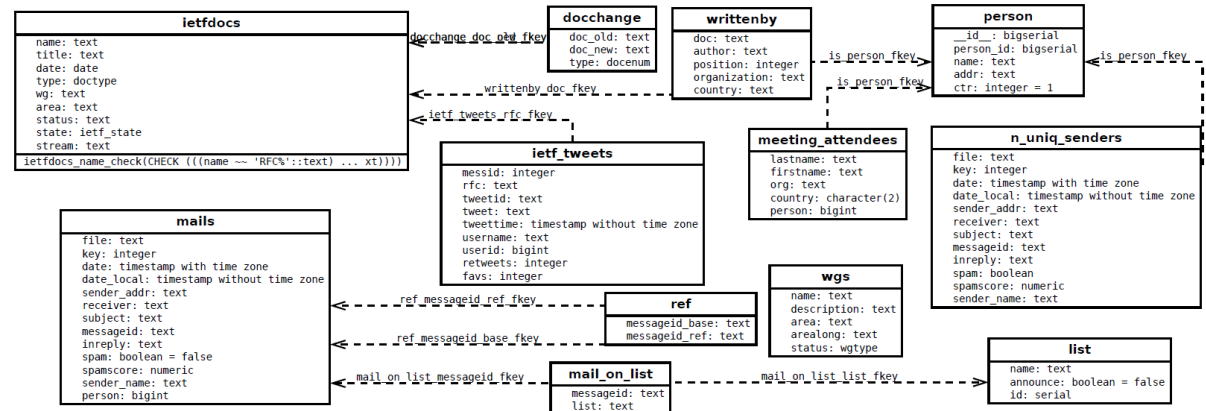
- Scope of work:
 - Understand how AC could be federated
 - Explore different AC mechanisms supporting federation, construct a taxonomy
 - Explore existing AC frameworks, enhance them to support federation

Text Mining on Mailing Lists: Sentiment Analysis

(Daniel, Heiko)

Database with IETF Data:

- Persons
- Documents
- Mails
- Working groups



Hypothesis: Individuals fulfill different roles / stereotypes in the discussions...

Your Task:

- Add sentiment information to the mails in the database
 - Mail content is not in the database but scripts to parse it exist
 - Evaluate different sentiment analysis tools

Practical topic!

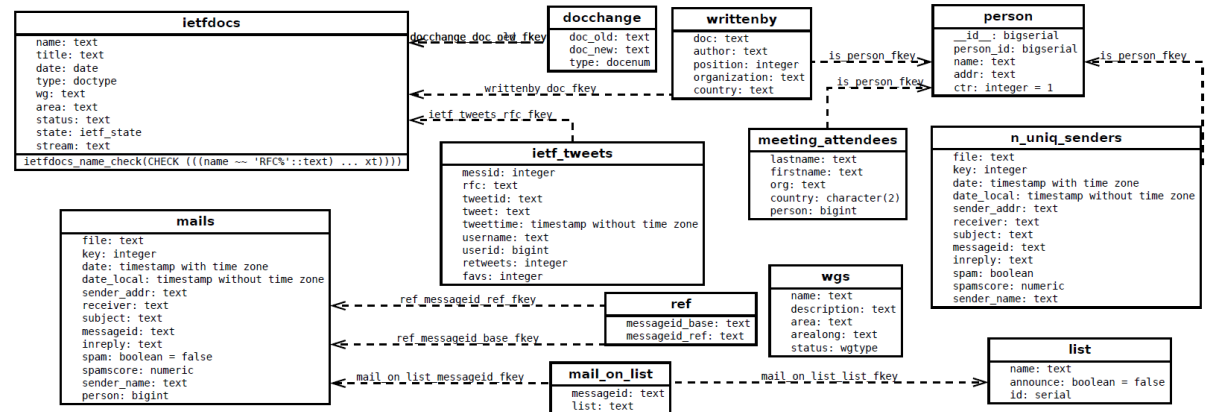
- Minimum page requirement is relaxed to 4 pages;
- Group work together with the second topic possible (“Text Mining on Mailing Lists: Tagging”)

Text Mining on Mailing Lists: Tagging

(Daniel, Heiko)

Database with IETF Data:

- Persons
- Documents
- Mails
- Working groups



Your Task:

- Add tags to mails and cluster content
 - Mail content is not in the database but scripts to parse it exist
 - Tag-based clustering is a simple way to cluster content
- Cluster the topics in the IETF

Practical topic!

- Minimum page requirement is relaxed to 4 pages;
- Group work together with the second topic possible (“Text Mining on Mailing Lists: Sentiment Analysis”)

Deterministic TCP Packet Sampling (Johannes)

Bidirectional TCP connections can reveal (?) a lot about the network:

- Latency in the network
- Latency of the application

But we are not always able/do not want to process all packets for performance monitoring. So we have to work with suitable subsets, i.e. not all packets, but we want bidirectional packet pairs for the ACK.

You Task:

- There is a very basic concept for such a sampling process, you should investigate its feasibility and explore the concept. What other work exists in the space already?
- Create a basic implementation (based on libpcap) and test it.

Requirement: Talk to me before choosing the topic!