

Using Networks to Teach About Networks (Report on Dagstuhl Seminar #17112)

Jürgen Schönwälder
Jacobs University
Bremen, DE
j.schoenwaelder@jacobs-university.de

Timur Friedman
UPMC Sorbonne Universités
Paris, FR
timur.friedman@upmc.fr

Aiko Pras
University of Twente
Enschede, NL
a.pras@utwente.nl

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

This report summarizes a two and a half days Dagstuhl seminar on “Using Networks to Teach About Networks”. The seminar brought together people with mixed backgrounds in order to exchange experiences gained with different approaches to teach computer networking. Despite the obvious question of *what to teach*, special attention was given to the questions of *how to teach* and *which tools and infrastructures can be used effectively* today for teaching purposes.

Keywords

Computer networks, Education

1. INTRODUCTION

Computer networks have become a common utility and the Internet provides new opportunities for education. In addition, we see an increasing deviation of the deployed Internet from the basic principles driving the design of computer networks. All this has an impact on how we educate young minds in computer networking and hence it is required to rethink how education in computer networking should be organized, which topics are essential to cover and which ones are merely nice illustrations of core concepts. Furthermore, it seems necessary to think about using the Internet itself more intensively to develop new educational materials.

In order to start a discussion of such educational aspects, a Dagstuhl seminar titled [Using Networks to Teach About Networks](#) has been organized. The seminar took place in March 12–15, 2017. Some questions discussed during the seminar were:

- Which topics should be taught in a typical undergraduate course? What are the essential basic principles that need to be understood? Which topics should be covered in a typical graduate course? How to deal with the fact that architectural concepts are often violated in real networks?
- How should topics be taught? How to best use the Internet for teaching how the Internet works? How can we more easily ‘mesh’ teaching materials? Can we better organize the sharing of video content, assignments, or experimental setups? Do we need an open source platform for teaching materials? What about open source books on computer networks replacing traditional textbooks?

- What is the experience with modern teaching styles, such as pure online courses like MOOCs or flipped classrooms? Which role should project work play? How can novel teaching ideas best be leveraged and integrated into existing educational concepts?

The seminar was attended by 26 people (a complete list can be found at the end of this report), most of them affiliated with academic institutions. This report summarizes some of the presentations and demonstrations given during the seminar and it highlights some ideas that have been discussed in breakout groups. We hope that this report encourages instructors of computer networking courses to try out new ideas and approaches and that it stimulates further discussions about how to best educate young minds about computer networking.

In Section 2, we summarize presentations given during the seminar. During the seminar, four ad-hoc breakout groups were established and we present their results in Section 3. Section 4 briefly describes demonstrations that were given during the seminar. We conclude the report in Section 5.

2. PRESENTATIONS

Several prepared presentations were given during the seminar. The slides of the presentations can be found on the shared documents page of the seminar [1].

2.1 Collaborative teaching and learning

Jordi Domingo-Pascual (UPC) discussed in his presentation which concepts to teach and at which level. He stressed the point that the real Internet can be used for teaching purposes and he further developed the idea of collaborative teaching, i.e., the option to run labs concurrently at multiple institutions and to let students collaborate over the Internet to do experiments with the Internet.

2.2 Anytime and anyplace learning

John Domingue (OU) stressed the need to support anytime and anyplace learning. He reviewed in his presentation how the EU-funded [FORGE](#) project has been providing tools that integrate network experimentation facilities developed by the [FIRE](#) project into an online learning system.

2.3 Active learning experience

Gunnar Karlsson (KTH) explained how he has redesigned his introductory computer networking course to move away

from teacher centered instruction towards active learning [10]. Active learning has been shown to increase student performance in science, engineering, and mathematics [9]. Gunnar redesigned his course by reducing the scope of what he teaches and leaving data communication as well as network architecture and standardization as self-study for the students. The course now has continuous examination in the form of five mini-exams, three mandatory laboratory sessions, two mandatory individual written reports and four mandatory case studies as group work with reports and presentations in class. Active learning is realized by posing a problem and letting students discuss solutions in smaller groups (2-3 students) before groups report their results and compiling a joint solution at the board.

2.4 Experience with the rake philosophy

Jean-Yves Le Boudec (EPFL) discussed that he sees two different options, namely to either teach all the details of all networking protocols (largely infeasible) or to be focused on the general principles, leaving the mountains of details to further study. Jean-Yves Le Boudec adopted the rake philosophy where he is covering depths by carefully selected labs and breadth by extrapolation based on lectures and labs. During classes, he uses an active learning approach where students are asked in a first step to invent their own solution to a given problem and in a second step the student's solution are compared to existing solutions. The idea is that students only have to learn the difference to their own solution.

2.5 Educating future systems programmers

Lisa Yan (Stanford) reported about their experiences with running an undergraduate networking course that stresses implementation work. The course material is centered on questions such as "How does the Internet work?", "What is the theory behind how the Internet works?", and "Why was the Internet developed this way?". Students spend a large amount of their time on intensive programming tasks in which students basically implement their own IP router from scratch. Tools like [Mininet](#), [VirtualBox](#), [Wireshark](#), and [Mahimahi](#) are used within an [OpenEDX](#) environment. The instructors use the [flipped classroom](#) approach in class meetings.

2.6 Educating future researchers

Lisa Yan (Stanford) reported that their graduate course is largely focused on reproducing research. Students first summarize a research paper and afterwards they try to reproduce the research results. Students are encouraged to interact and collaborate with other researchers, in particular the authors of the original research papers the students are trying to reproduce. Letting graduate students reproduce research has been found beneficial for the students since they have to understand a paper in detail and they build up a personal relationship with the authors. Furthermore, the knowledge that a research results has been reproduced is valuable for the research community as a whole.

2.7 Using learning analytics

Marc-Oliver Pahl (TUM) talked about [learning analytics](#), i.e., the measurement, collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimizing learning and the environments

in which it occurs. Marc-Oliver Pahl is using learning analytics intensively in his courses and labs to continuously improve teaching. Students can always see their results and their relative ranking. He recently started experiments trying to predicting learning outcomes. This, of course, can also be risky as such predictions may change the student's attitude towards a course or lab.

2.8 Recording learning achievements

John Domingue (OU) discussed the usage of [blockchain](#) technology in order to record learning achievements. The basic idea is to move the storage of data about achievements from organizations issuing certificates to a distributed blockchain. The benefit is that data is owned and controlled by students instead of educational institutions while increasing transparency and reducing risks of fraud.

3. BREAKOUT GROUPS

The seminar participants did split into four smaller groups in order to discuss some topics in more detail. The following sections summarize the results of the breakout group discussions.

3.1 Content of computer networking courses

Most people follow, at least partially, traditional textbooks (e.g., the James Kurose and Keith Ross textbook [11]) or online textbooks such as Olivier Bonaventure's computer networking book [5] or Jean-Yves Le Boudec's tutorial on rate adaptation, congestion control, and fairness [6]. While there is a common core of topics that people seem to cover, there are also many differences due to the different functions courses have in the various curricula or differences in the target audiences. Topics typically covered are Internet architecture, physical layer, link layer, IP layer, intra-domain routing, inter-domain routing, congestion control, application layer protocols, network security, building simple networks, practical assignments (a more detailed discussion can be found on the shared documents page of the seminar [1]). There are, however, often significant differences in the details and in which order and depths topics are covered.

Overall, it seems desirable to move towards a modular framework of composable educational units. Such a framework could facilitate the exchange and evolution of educational material. Educational material in this context covers textual resources (books, book chapters, articles, ...), presentation slides, videos, exercise sheets, programming tasks, lab experiments, quizzes, and exam questions. Furthermore, it is desirable to add metadata to educational material, such as authors, editors, contributors, license conditions (preferably creative commons BY). In addition, it seems to be useful to track where educational materials are used. It was also suggested to discuss course units in the context of competence levels, for example based on the Revised Bloom Taxonomy [12], which distinguishes in the cognitive dimension *remembering*, *understanding*, *applying*, *analyzing*, *evaluating*, and *creating* and in the knowledge dimension *factual knowledge*, *conceptual knowledge*, *procedural knowledge*, and *metacognitive knowledge*.

While there was general agreement that it is desirable to more easily share educational materials, it is less clear how to reward people for sharing material in a form that is easily reusable and which kind of infrastructure is necessary to organize the sharing process.

3.2 Teaching methods

Computer networking courses differ based on the target groups (primarily electrical engineering students versus primarily computer science students) and on the place of the course in the curriculum (typically 2nd or 3rd year). In general, students tend to have problems switching between different views and facets of a concept (e.g., understanding the interplay of different protocol layers) and thinking in terms of an asynchronous distributed system.

A general goal of teaching methods is to encourage students to be active, i.e., to make them ask questions or to let them develop solutions to certain computer networking problems. The following teaching methods were discussed in more detail:

- **Flipped classroom:** The flipped classroom teaching method assumes that teaching materials are studied by students at home before the class session, while in-class time is devoted to discussions or exercises [7, 8]. A common problem is that students who are not used to flipped classroom style of teaching often come unprepared or they misunderstood the content. This makes discussion in classes sometimes difficult and courses can mutate into “sandwich classrooms”, where students first self-study before having a classroom discussion, often followed by another self-study phase.
- **Students grading students:** Several people reported positive experience with letting students grade results produced by other students or letting students create teaching materials that are reviewed by their peers. Overall, students tend to be fair, they often grade tougher than the regular instructors. Of course, involving students not only in the production of solutions to given problems but also in the assessment of solutions is not a cheap grading tactic; instructors need to carefully monitor the process and they are in charge of grading the student’s assessments. Using the students grading students approach requires that clear guidelines are provided, that expectations are clearly communicated, and that sufficient anonymity is provided (double blind), which may require a certain minimum class size. A radical example of this educational approach is [École 42](#) and [42 USA](#).
- **Conference-style seminars:** Some attendees reported about their positive experience with letting students write reports about selected topics and to organize a double-blind review process where students have to evaluate reports written by other students. Students are allowed to revise their reports based on the reviews before giving a short presentation about the topic in class. Grading is based on the reviews a student has written and the presentation, but not on the report itself.
- **Student competitions:** Some attendees reported positive experiences with posting challenges that lead to competitions between student groups. The challenges are well defined tasks that must be solved in a given timeframe. In order to stimulate competition, it is crucial to have a live scoring system providing student teams with direct feedback about their performance relative to others. Grading depends on the

achievement of the student teams. It is possible to include a presentation of the winning solution at the end. Student competitions require that a longer timeslot is available, ideally a day or at least half a day, so that students can concentrate on the task given to them.

The sizes of computer networking courses vary significantly between different academic institutions. Scaling courses to large numbers of students requires careful planning, in particular when it comes to lab sessions or programming assignments. It is important to find ways to prevent plagiarism. For program code, structural similarity testing tools like [MOSS](#) can be useful. Systems like [Turnitin](#) can help detect plagiarism in written reports. It is important that any usage of such tools is announced well ahead of the assignments, ideally at the beginning of a course or lab. For communication outside the classroom, collaborations systems like [Piazza](#) have been found useful. Some institutions use fully fledged online learning platforms such as [OpenEDX](#) or [Moodle](#).

3.3 Tools and testbeds

In addition to regular command line tools, a number of more specialized tools are already widely used in lab courses and to a lesser extent in classrooms. [Wireshark](#) is widely used to dissect packets and to analyze captured packet traces. Wireshark is also good for understanding packet flows or specific protocol interactions. There are also some repositories of open packet traces ([2–4]) that can be used in student projects. Commonly used flow analysis tools are [Bro](#), [Tranalyzer](#), [ntop](#), or [nfdump](#). A powerful packet generation tool is [Scapy](#).

Network emulation tools seem to be replacing network simulators such as [NS3](#). Emulation tools are able to scale up to the sizes typically needed for undergraduate courses (or labs) and the learning curve is usually lower. [Mininet](#) seems to be a popular solution together with its cousin [Mininetx](#), which however does not seem to be actively maintained. There are in addition graphical network emulation tools such as [GNS3](#), which can also emulate command line interfaces of commercial routers.

Different approaches can be used to make experiments on the Internet. The [PlanetLab](#) platform can be used to let students gain experience with running software on a live distributed system. However, for simple experiments, it has been found useful to implement a more student friendly interface on the PlanetLab infrastructure that makes it easy to run simple experiments without all the usual PlanetLab account and slice management overhead. This approach has been used in Timur Friedman’s network measurement MOOC. The [RIPE Atlas](#) measurement infrastructure has been found easy to use for network measurements, in particular due to the availability of easy to use APIs. The same is true for the [RIPE Stat](#) service, which makes it easy to retrieve a lot of metadata about the Internet resources, both via a web interface or via an API.

For many labs, it is useful to have access to good visualization tools. It is a benefit if students already know standard tools like [gnuplot](#) or statistical analysis tools like [R](#). Some specific visualization tools that have been found useful are [BGPlay](#) and [TPlay](#). Visualization tools that have been found missing are tools that can properly visualize network dynamics.

3.4 Educational technology

Educational technology can be used to scale up courses to large numbers of students or to allow students to study at their own pace independent of classroom meetings that are imposing a fixed learning pace on all students. Furthermore, educational technology can deliver detailed insights about how students learn and which topics they find difficult to understand. Typical problems that were experienced while using educational technology are related to cheating, keeping students motivated, and keeping students focused. Technical setup problems still exist although things seem to improve. Problematic are tutoring interactions (many questions pop up during the night before a deadline) and there is generally a lack of useful and actionable feedback.

Cheating problems can be reduced by having a strong authentication system (Coursera for example uses fingerprints and webcam pictures). Hardware authentication devices such as [YubiKey](#) may further help in multi-factor authentication systems. Another helpful approach to reduce cheating is randomization or even personalization of tests.

In order to keep students motivated, it is useful to present content in small units and to integrate questions regularly to assess the learning progress. It is also useful to construct breaks by switching learning media frequently, e.g., switching from video content to a quiz, then back to video content followed by a practical experiment and so on. Another motivator can be some form of competitions. It can be useful to think of a course as a computer game with multiple levels that can be reached.

The goal of learning analytics is generally to improve learning materials and keeping in touch with the virtual learner groups of an online course. Online learning systems allow to collect a lot of data but it remains unclear which information should be collected and which information should not be collected. There are certain ethical and legal considerations and of course privacy concerns. For example, should the time spent per learning element be used to customize tests or exams? How comparable are such personalized exams? What about correlation with demographic data? And who (student, tutor, instructor) is allowed to have access to which data (and for which purposes)?

Since the production of online learning material is very time intensive and hence expensive, it is useful to find ways to collaborate and to share learning materials. However, there is also a value of a diversity of teaching approaches. By sharing educational units at a large scale, there is a certain risk that less people will be thinking about how to best explain certain concepts and hence we may lose valuable fresh ideas.

4. DEMONSTRATIONS

Seminar participants were invited to demonstrate educational approaches or specific tools that they found useful. The following sections summarize some of the demonstrations given during the seminar.

4.1 Blended learning for teaching networks

Marc-Oliver Pahl (TUM) demonstrated iLab, a blended learning hands-on course concept. The didactic concept builds on four phases: (a) lecture (1.5 hours), (b) individual preparation (≈ 6 hours), (c) practical teamwork (≈ 10 hours), and (d) individual oral exams (20 minutes). An eLearning system has been implemented to support these

four phases and to collect data for learning analytics. The didactic approach has been used successfully with approximately 2000 Bachelor and Master students so far.

4.2 Internet security MOOC

Aiko Pras (UT) and Anna Sperotto (UT) showed their work on a [MOOC on Internet Security](#), running on the [OpenEDX](#) platform. They have created short explanatory videos and student exercises that are often customized for each student. For example, they create different traffic traces for each student, which makes it difficult for students to simply copy a solution created by some other student. The Internet security MOOC is currently running at a small scale for testing purposes. Aiko Pras reported that the availability of the OpenEDX infrastructure at the University of Twente already motivated other colleagues to use the online learning infrastructure for their courses as well. Hence, you will find a collection of additional courses on the platform that were not initially envisaged.

4.3 Student competitions

Pieter-Tjerk de Boer (UT) is successfully using student competitions for educational purposes and he intends to make them available to other universities. He demonstrated an assignment where students have to design and implement a medium access control protocol to share a time-slotted medium fairly and efficiently among four nodes. The students are provided with a template of a program that decides whether a node uses an announced time-slot or not. The challenge given to the students is to design algorithms that try to avoid collisions and improve the utilization of the channel. Student groups design and implement their own algorithms and run them against a server. The server calculates efficiency and fairness scores that are immediately shown to all participants. This immediate feedback motivates students to engage in a competition between student groups, which generally improves student activity and learning outcomes.

4.4 Measurement data analysis exercise

Fabio Ricciato (UL) explained how he is teaching the pitfalls of measurement data analysis. He provides students with datasets together with a short description how the datasets have been produced. The students are given the task to analyze the dataset and answer some (apparently) simple questions. The assignment is inspired by common problems that are typically encountered in real dataset, such as incomplete context information and ambiguous meta-data, and it is designed to expose the risks of a superficial (mis)use of the most basic statistical concepts. Fabio Ricciato did run his toy measurement data analysis exercise as part of the PhD school on traffic monitoring and analysis, which was part of [8th Traffic Monitoring and Analysis](#) workshop (TMA 2016). In general, letting students make mistakes they can learn from seems to be an effective approach. Another useful method is to review mistakes made by others, e.g., by critically discussing with the students the methodological pitfalls encountered in some papers. The key message made by Fabio Ricciato was that inducing the students to discover “how NOT to do things” is not less important than explaining directly how things should be done — a pedagogical attitude that Fabio Ricciato summarizes as “teaching by negatives”.



4.5 Traffic mining and analysis

Stefan Burschka (RUAG) provided an overview how he is teaching traffic mining and troubleshooting techniques. His approach is to confront students with scenarios where it is necessary to develop creative approaches to solve puzzles given to students. The idea is to motivate students to pay attention to little details while at the same time students should learn that data always exists in a certain context that is very important in order to interpret the data in a correct way. In order to mine large datasets (packet captures larger than 1 TB) effectively, he is developing an extensible tool called *Tranalyzer*, that can efficiently extract flows without being bound to a very narrow definition of a traffic flow. Stefan Burschka did run his traffic mining exercise as lab sessions of the 10th *Autonomous Infrastructure, Management and Security* conference (AIMS 2016).

5. CONCLUSIONS

It became clear during the seminar that the way people teach computer networking courses is undergoing changes. During the time available at the seminar, it was possible to establish a common sense about the various ideas tried at different institutions. A reoccurring topic are the high costs for the production and maintenance of educational materials. In particular, the production and maintenance of good laboratory assignments is very time intensive. It would be nice if there were ways to organize more effective collaboration in order to more easily share educational materials and to mesh course and lab components.

Acknowledgments

This seminar was located at the International Conference and Research Center for Computer Science Schloß Dagstuhl in Wadern, Germany, supported by German federal and state funds. The organizers would like to thank the participants (alphabetically ordered by last name) for their contributions: *Olivier Bonaventure* (UC Louvain – Louvain-la-Neuve, BE), *Alessio Botta* (University of Naples, IT), *Stefan Burschka* (RUAG – Bern, CH), *Georg Carle* (TU München, DE), *Marinos Charalambides* (University College London, GB), *Pieter-Tjerk de Boer* (University of Twente – Enschede, NL), *Wouter de Vries* (University of Twente – Enschede, NL), *Jordi Domingo-Pascual* (UPC – Barcelona, ES), *John Domingue* (The Open University – Milton Keynes, GB), *Timur Friedman* (UPMC Sorbonne Université – Paris,

FR), *Gunnar Karlsson* (KTH Royal Institute of Technology – Stockholm, SE), *Jean-Yves Le Boudec* (EPFL – Lausanne, CH), *Simone Opel* (Berufskolleg der Stadt Bottrop, DE), *Marc-Oliver Pahl* (TU München, DE), *Cristel Pelsser* (University of Strasbourg, FR), *Aiko Pras* (University of Twente – Enschede, NL), *Bruno Quoitin* (University of Mons, BE), *Fabio Ricciato* (University of Ljubljana, SI), *Jürgen Schönwälder* (Jacobs University – Bremen, DE), *Anna Sperotto* (University of Twente – Enschede, NL), *James P. G. Sterbenz* (University of Kansas – Lawrence, US), *Burkhard Stiller* (Universität Zürich, CH), *Iyad Tumar* (Birzeit University, PS), *Klaus Wehrle* (RWTH Aachen University, DE), Harald Weikert (IsarNet AG – München, DE), *Lisa Yan* (Stanford University, US). This preparation of the seminar was funded by Flamingo, a Network of Excellence project (ICT- 318488) supported by the European Commission under its Seventh Framework Programme.

6. REFERENCES

- [1] Dagstuhl Seminar #17112 – Using Networks to Teach About Networks: Materials. <http://materials.dagstuhl.de/index.php?semnr=17112>. [Online; accessed 08-June-2017].
- [2] NETRESEC Pcap Files. <http://www.netresec.com/?page=PcapFiles>. [Online; accessed 09-June-2017].
- [3] Simple Web Traces. <https://traces.simpleweb.org/>. [Online; accessed 09-June-2017].
- [4] Wireshark Sample Captures. <https://wiki.wireshark.org/SampleCaptures>. [Online; accessed 09-June-2017].
- [5] O. Bonaventure. *Computer Networking: Principles, Protocols and Practice*. <http://cnp3book.info.ucl.ac.be/>, 2 edition, 2016.
- [6] J.-Y. L. Boudec. *Rate adaptation, Congestion Control and Fairness: A Tutorial*. http://ica1www.epfl.ch/PS_files/LEB3132.pdf, 2016.
- [7] C. H. Crouch and E. Mazur. Peer instruction: Ten years of experience and results. *American Journal of Physics*, 69(9):970–977, 2001.
- [8] L. Deslauriers, E. Schelew, and C. Wieman. Improved learning in a large-enrollment physics class. *Science*, 332(6031):862–864, 2011.
- [9] S. Freeman, S. Eddy, M. McDonough, M. Smith, N. Okoroafor, H. Jordt, and M. Wenderoth. Active learning increases student performance in science, engineering, and mathematics. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 111(23):8410–8415, 2014.
- [10] G. Karlsson and S. Janson. The flipped classroom: a model for active student learning. In L. Engwall, E. Corte, and U. Teichler, editors, *From books to MOOCs? Emerging models of learning and teaching in higher education*, pages 127–136. Portland Press Limited, 2016.
- [11] J. F. Kurose and K. W. Ross. *Computer Networking: A Top-Down Approach*. Pearson, 6 edition, 2012.
- [12] L. O. Wilson. *Anderson and Krathwohl – Bloom’s Taxonomy Revised*. <http://thesecondprinciple.com/teaching-essentials/beyond-bloom-cognitive-taxonomy-revised/>, 2016.