

Report on the Dagstuhl Seminar on Visualization and Monitoring of Network Traffic

**Daniel A. Keim · Aiko Pras · Jürgen Schönwälder ·
Pak Chung Wong · Florian Mansmann**

1 Introduction

The Dagstuhl Seminar on Visualization and Monitoring of Network Traffic [1] took place May 17–20, 2009 in Dagstuhl, Germany. Dagstuhl seminars promote personal interaction and open discussion of results as well as new ideas. Unlike at most conferences, the focus is not solely on the presentation of established results but also, and in equal parts, to presentation of results, ideas, sketches, and open problems. The aim of this particular seminar was to bring together experts from the information visualization community and the networking community in order to discuss the state of the art of monitoring and visualization of network traffic. People from the different research communities involved jointly organized the seminar. The co-chairs of the seminar from the networking community were Aiko Pras (University of Twente) and Jürgen Schönwälder (Jacobs University Bremen). The co-chairs from the visualization community were Daniel A. Keim (University of Konstanz) and Pak Chung Wong (Pacific Northwest National Laboratory). Florian Mansmann (University of Konstanz) helped with producing this report. The seminar was organized and supported by Schloss Dagstuhl and the European Network of Excellence for the Management of Internet Technologies and Complex Systems (EMANICS) [2].

D. A. Keim · F. Mansmann
University of Konstanz, Constance, Germany

A. Pras
University of Twente, Enschede, The Netherlands

J. Schönwälder (✉)
Jacobs University Bremen, Bremen, Germany
e-mail: j.schoenwaelder@jacobs-university.de

P. C. Wong
Pacific Northwest National Laboratory, Richland, WA, USA

2 Motivation

The seamless operation of the Internet requires being able to monitor and to visualize the actual behavior of the network. Today, IP network operators usually collect network flow statistics from critical points of their network infrastructure. Flows aggregate packets that share common properties. Flow records are stored and analyzed to extract accounting information and increasingly to identify and isolate network problems or security incidents. Whereas network problems or attacks that significantly change traffic patterns are relatively easy to identify, it tends to be much more challenging to identify creeping changes or attacks and faults that manifest themselves only by very careful analysis of initially seemingly unrelated traffic patterns and their changes. There are currently no deployable good network visualization solutions supporting this kind of network analysis, and research in this area is just starting. In addition, the large volume of flow data on high capacity networks and exchange points requires moving to probabilistic sampling techniques, which require new analysis techniques to calculate and also to visualize the uncertainty attached to data sets.

3 Seminar Scope

The aim of the seminar was to bring together for the first time people from the networking community and the visualization community in order to explore common grounds in capturing and visualizing network behavior and to exchange upcoming requirements and novel techniques. The seminar also targeted network operators running large IP networks as well as companies building software products for network monitoring and visualization. We believe that bringing together experts from two usually separate fields made this seminar unique and we expect that the intensive exchange in a Dagstuhl seminar setting has high potential to lead to joint follow-up research. The following research questions were suggested for discussion:

1. What are suitable data analysis and visualization techniques that can operate in real-time and support interactive online operation?
2. How can monitoring and visualization techniques be made scalable?
3. How can distributed monitoring systems be self-organizing and adapt dynamically to changes in network and service usage?
4. How can algorithms aggregate data within the network and trade accuracy of the measurement results against data collection overhead?
5. What are suitable sampling techniques and how does sampled data impact data analysis techniques and data visualization?
6. Which filtering, zooming, and correlation techniques can be applied in real-time?
7. What are good techniques for visualizing unusual traffic patterns or very rare patterns?
8. What are effective methods to detect and to visualize intrusions, like (distributed) scan attempts and denial of service attacks?

While this list was helpful as an orientation, not all of the research questions were discussed at the same depth during the seminar. Moreover, other concerns, such as the efficient storage and processing of large amounts of NetFlow/IPFIX records, were emphasized in the presentations and discussions.

4 Program

The first day started with three keynote talks in order to set the stage. The following keynotes were presented:

- Network Visualization (Jack van Wijk, Eindhoven University of Technology, The Netherlands)
- Network Traffic Visualization with IsarFlow (Harald Weikert, IsarNet, Germany)
- On Detection and Analysis of BGP Anomalous Dynamics (Felix Wu, University of California at Davis, USA)

The network visualization keynote by Jack van Wijk was an introduction into graph drawing, information visualization for the participants from the network research community. The talk not only introduced basic concepts such as graph drawing, information visualization (the visualization pipeline and a set of basic visualizations), and visual analytics, but also highlighted some properties of the human visual perception.

Harald Weikert discussed in his keynote the visualizations currently available in commercial network management products and how they address the requirements of network operators his company works with. He raised the fundamental question whether the direction is to enhance visualizations into more complex charts or to enhance the processing of the data in order to show only relevant information in easy to comprehend charts. While there are tradeoffs between these two directions, participants agreed that research in both directions is needed.

The third keynote given by Felix Wu focused on the analysis of Internet routing anomalies, a time intensive process requiring highly skilled people. By going through some example studies of the past, he came to the observation that the main challenge is that there is on the one hand too much information while analyst often do not have the information needed to fully understand an observation. Things are often further complicated by a lack meta data describing how information was generated. He concluded that visualizations should be strongly linked to explanations to be useful.

In the afternoon, participants had the opportunity to give short talks presenting their research activities in the fields of networking and visualization. Twenty-three participants made use of this opportunity, leading to a very interesting session exploring the seminar topic from very different viewpoints. The evening session was reserved for demonstrations of research prototypes. Several people brought their software to the meeting and enjoyed showing their programs to peers and for many attendees the demonstration evening became one of the highlights of the seminar.

The second day started by forming four working groups in order to inspire direct interactions and fresh and novel research. Thereby, particular care was taken that the groups were mixed in a way that each one contained experts of both fields. Each of the following questions was assigned to two of the four working groups.

1. What could be applications of visualization in the area of network monitoring?
2. What visualization could be useful to solve networking problems?

After the discussions within the working groups, the results were shared and discussed with all participants of the seminar.

The last day was used to continue discussions in a plenary meeting. Some topics for publications were identified that could help people to get better involved in network data visualization. In particular, a survey of data visualization techniques that have been applied to visualize network traffic would be useful to have. Furthermore, it would be nice to have a document describing requirements for visualizations coming from the networking community. Similarly, a paper from the visualization community describing visualization techniques and associated case studies where visualization techniques have been used successfully to solve real-world problems would be useful. Finally, papers discussing experiences from using visualization techniques in operational environments reflecting on lessons learned and outlining open research issues could be very helpful.

5 Conclusions

The seminar was a fertile meeting in which 36 researchers with diverse background met.

While most attendees came from Europe, the seminar enjoyed a significant number of participants from the USA as well as participants from South Korea, Australia, New Zealand and Brazil. Most participants were affiliated with universities or state-owned research centers while some participants were employed by industry or industrial research centers. The diversity of the background of the participants resulted in interesting and useful discussions, new understandings of the fundamental concepts and problems in the field, and in new collaborations on an array of problems which were not well defined or identified prior to this seminar.

Several work groups created during the seminar not only generated new insights into specific topics in the field of visual network monitoring, but also initiated ongoing joint work, with group members continuing the work they started at the seminar. The seminar included multiple presentations and discussions. In particular, the largely disjoint research communities of Networking and Visualization exchanged their methods and unsolved problems, resulting in fruitful discussions and awareness of the other field. This seminar clearly illustrated the diversity, relevance, and fertility of the topics we presented and discussed. The intensity of the participants' involvement leads us to believe that the interactions fostered by the seminar will generate a lot of follow-up research, and eventually lead to practical use as well.

References

1. Seminar web page: <http://www.dagstuhl.de/09211/>. Last accessed Nov 2009
2. EMANICS web site: <http://www.emanics.org/>. Last accessed Nov 2009

Author Biographies

Daniel A. Keim received a PhD degree in computer science from the University of Munich in 1994. He is a full professor in the Computer and Information Science Department, University of Konstanz. Dr. Keim was program co-chair of the IEEE Information Visualization Symposia in 1999 and 2000, the ACM SIGKDD Conference in 2002, the VisSym Conference in 2004, and the Visual Analytics Symposium in 2006. Currently, he is on the editorial board of the IEEE Transactions on Knowledge and Data Engineering, the Knowledge and Information System Journal, and the Information Visualization Journal.

Aiko Pras is Associate Professor at the Design and Analysis of Communication Systems (DACS) group at the University of Twente (UT), the Netherlands. His research interests include network management technologies, Web services, network measurements, and Internet security. He chairs the IFIP Working Group 6.6 on Management of Networks and Distributed Systems. He was the technical program co-chair of the Ninth IFIP/IEEE Integrated Management Symposium (IM 2005), has been a Steering Committee member of the IFIP/IEEE NOMS and IM Symposia (NISC), and was general co-chair of Manweek 2009.

Jürgen Schönwälder is Associate Professor of Computer Science at Jacobs University Bremen, Germany. He received his doctoral degree in 1996 from the Technical University Braunschweig, Germany. His research interests are distributed systems, network management, wireless sensor networks, and network security. He is an active member of the Internet Engineering Task Force (IETF) and chair of the Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF). He serves on the editorial board of the IEEE Transactions on Network and Service Management and the Springer Journal of Network and Systems Management.

Pak Chung Wong is a chief scientist and project manager at the Pacific Northwest National Laboratory in Richland, Washington, USA, where he performs research and development on information technology and scientific computation. His research interests include visual analytics, visualization, social computing, bioinformatics, human-computer interaction, privacy and security, and computational science. He received a Ph.D. in computer science from the University of New Hampshire.

Florian Mansmann is a research scientist and lecturer at the University of Konstanz in Germany, where he obtained his Ph.D. about Visual Analytics in the field of Network Monitoring and Security. His main research interests are in Visual Analytics with a focus on its application to Network Monitoring and Security, Data Mining, Spatiotemporal Data Analysis, and Sentiment Analysis. Mr. Mansmann is a program committee member of the VizSec workshop on Visualization for Cyber Security.