

09211 Abstracts Collection
Visualization and Monitoring of Network Traffic
— Dagstuhl Seminar —

Daniel A. Keim¹, Pak Chung Wong², Aiko Pras³ and Jürgen Schönwälder⁴

¹ Universität Konstanz, D
keim@uni-konstanz.de

² Pacific Northwest National Laboratory, Richland, USA
pak.wong@pnl.gov

³ University of Twente, NL
pras@cs.utwente.nl

⁴ Jacobs University - Bremen, D

Abstract. From 17.05. to 20.05.2009, the Dagstuhl Seminar 09211 “Visualization and Monitoring of Network Traffic ” was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Computer Networks, Internet, Monitoring of Networks and Services, Visualization Animation

09211 Executive Summary – Visualization and Monitoring of Network Traffic

The seamless operation of the Internet requires being able to monitor and visualize the actual behaviour 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. While network problems or attacks significantly changing 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 pattern and their changes. There are currently no deployable good solutions and research in this area is just starting. In addition, the large volume of flow data on high capacity networks and exchange points requires to move to probabilistic sampling techniques, which require new analysis techniques to calculate and also visualize the uncertainty attached to data sets.

Keywords: Computer Networks, Internet, Monitoring of Networks and Services, Visualization Animation

Joint work of: Keim, Daniel A.; Pras, Aiko; Schönwälder, Jürgen; Wong, Pak Chung; Mansmann, Florian

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2009/2157>

The Usability of Information Visualisation Techniques for Network Monitoring

Keith Andrews (TU Graz, AT)

The field of information visualisation has given birth to a wide palette of interactive techniques to visualise a variety of types of abstract information. The field of network traffic monitoring generates a huge amount of data and information of varying types.

The question now arises: which visualisation technique best supports a network analyst for which common tasks? And how can a suite of techniques be best combined into a consistent visual interface for network monitoring?

The answer: a heady mixture of network data harvesting, filtering, transformation, information visualisation, user-centered design, and usability testing.

Keywords: Information visualisation, user-centered design, usability testing.

Working Group Results "We have a hammer, find a nail"

Stephan Diehl (Universität Trier, DE)

Looking at various visualization techniques we discussed how to use them for network management.

High-speed Monitoring and Intelligent data pre-selection for Attack Detection

Falko Dressler (Universität Erlangen, DE)

Monitoring in high-speed networks requires novel solutions with respect to packet selection and information processing. We present selected solutions based on flow monitoring and packet sampling mechanisms that allow to reduce the amount of monitoring data in a situation-aware manner. For example, we introduce Front Payload Analysis (FPA) that extends flow monitoring by selecting the first N payload bytes of each flow for signature based intrusion detection. Furthermore,

we discuss monitoring techniques for deep packet inspection. Basically, the administrators have to restrict monitoring to selected parts of traffic in case of potential overload. The main drawback is that either a high proportion of benign traffic is needlessly processed by the computationally intensive IDS, or a static configuration needlessly excludes traffic from intrusion detection. We developed a self-configuring solution that tries to alleviate both approaches by introducing a method for intelligent filtering of traffic data for IDS and dynamic host-based traffic selection based on pre-defined priorities and detected anomalous events. Both approaches decrease the amount of data to be processed at the IDS substantially, while trying to avoid a static configuration that enables attackers to anticipate monitoring holes.

Keywords: Network monitoring, adaptive reconfiguration, flow monitoring

Joint work of: Dressler, Falko; Limmer, Tobias

Cyber Analytics: Challenges and solutions for computer security

Glenn A. Fink (Pacific Northwest National Lab., US)

At Pacific Northwest National Laboratory, we are defining a new area of inquiry we are calling Cyber Analytics. Cyber analytics is observing computer and network data, and quantifiably comparing it to theoretical behavioral models to support decision-making. Informally the discipline is to understand the behavior of computers and computer networks from the data they generate. This presentation tells the needs that drive the creation of this discipline and outlines some of the solutions that are needed.

Keywords: Computer security, cyber analytics, analytics, visualization

Interactive Exploration of Typed Networks

Carsten Goerg (Georgia Institute of Technology, US)

Networks are often represented as node-link diagrams and visualized as a set of circles (or other shapes) and lines connecting them. However, there exist many other visual representations for networks that depict a set of entities and connections between them.

Different types of visualizations along with different interaction techniques support different types of tasks. Combining multiple visualizations to provide different perspective on the network can be even more powerful than a single one of them. In my presentation, I will walk through some examples to show which (combination of) visualizations are best suited for which type of task.

Accounting system for heterogeneous IP-networks (IPNA) implemented at Kaiserslautern University

Hans Hagen (TU Kaiserslautern, DE)

This paper describes an accounting system (IPNA) for heterogeneous IP-networks with arbitrary topologies implemented at the university of Kaiserslautern. The produced data volume per unit is numerated. The collected data is stored in a database and offers different analysis possibilities. The results can be visualized and adapted to the users requirements.

The main effort was to build a data traffic quota system for single units as well as groups of devices that also report exceeded quotas. The system itself only observes the network traffic. Interfaces offer tools to interact with the network. The IPNA consists of a back-end for the data-acquisition and -preparation and a front-end for configuration and visualization tasks including quality control.

Keywords: Accounting system, IP-network, Communication, information visualization, online quality control

Joint work of: Worden, Brian; Baltes, Claudia; Scheler, Inga; Müller, Paul; Hagen, Hans

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2009/2155>

Visualizations in Network Operations and Management: What Works, What Doesn't, and What's Missing

Simon Leinen (Switch - Zürich, CH)

This presentation looks at the use of data visualization at a research network operator over more than a decade. We will consider various audiences and uses for these visualizations, but primarily day-to-day use in network management, including monitoring, traffic planning, and troubleshooting. A particular focus is on visualizations that would seem useful (or at least interesting to try), but aren't yet available.

Keywords: Internet, network management, visualization, data mining

Cluster Visualization in Network Traffic

Lars Linsen (Jacobs University - Bremen, DE)

Network traffic has characteristics determining its behavior. In order to find patterns and outliers in network traffic we look into the multidimensional space formed by the parameters that describe the traffic's characteristics. We apply multidimensional data visualization consisting of an automated clustering step and an interactive exploration step.

The interactive exploration uses linked views, integrated views, and feature space views.

Keywords: Multidimensional cluster visualization

Monitoring and Intrusion Detection with NFlowVis

Florian Mansmann (Universität Konstanz, DE)

In this talk we present the network traffic visualization tool NFlowVis. Starting from Intrusion Detection alerts, the tool interactively guides the analyst to the home-centric view, which shows all connectivity information between the selected attackers from the ID alerts and the hosts of the monitored network prefix.

Keywords: Intrusion detection, visualization, netflows, NFlowVis

FloVis a visual paradigm for forensic network data analysis

John McHugh (Dalhousie University, CA)

While a substantial amount of visualization has been applied to network data, especially in support of security analysis, the utility of most systems is dubious. In our work, we are primarily interested in discovering flexible, interrelated visual paradigms that can aid the analysts in understanding previously unseen phenomena. Since network data is inherently multiresolution, we need representations that can cover volumes ranging from data crossing the borders of networks that account for multiple /8s down to individual host behaviors. In earlier work a visual representation showed an emergent network behavior that only manifest on a very large scale. The current effort allows multiple views that can cover a wide range of scales. The FloVis framework accommodates a series of plugins to allow drill down and pivoting as well as providing query access to the underlying NetFlow data.

Keywords: NetFlow, Bundle Diagrams, Activity diagrams, NetBytes Viewer, OverFlow

Joint work of: Taylor, Teryl; Paterson, Diana; Glanfield, Joel; Brooks, Stephen; Gates, Carrie; McHugh, John

Using Space-Filling Curves in Visualization of Network Traffic

Taghrid Samak (DePaul Univ. - Chicago, US)

Network monitors produce a huge amount of traffic data continuously.

This data needs to be stored and communicated for analysis purposes.

An efficient traffic representation is needed to ensure that the maximum information is passed using reasonable resources. We consider using space-filling curves (SFC) in traffic visualization. SFCs provide many desirable properties that help characterize traffic flows and identify anomalous behavior. First, we propose a methodology for representing traffic using SFCs. The proposed method generates traffic images that is robust against compression. Compressed images provide both storage and bandwidth savings, for example, when the image is transmitted from the monitor to an analyzing engine. Second, the proposed technique is evaluated for different SFCs by comparing original images with compressed images in terms of mapping accuracy and efficiency.

The resulting images are shown to withstand aggressive compression while preserving traffic properties.

Keywords: Network traffic, space-filling curves

Joint work of: Samak, Taghrid; Ghanem, Sahar; Ismail, Mohamed A.

Explorative Visualization of Log Data to support Signature Development and Forensic Analysis

Sebastian Schmerl (BTU Cottbus, DE)

Intrusion detection systems (IDS) have been proved to be an important mean for protecting systems. Most IDSs deployed today realized a signature-based detection. Here logged events are compared with defined patterns (signatures) that indicate security violations. False alerts are impossible by definition, but the reality shows a different picture. The reasons for this are imprecise signatures, which results from the complexity of the signature development. In particular the derivation of signatures from given exploits are difficult. Manual log analysis is the basis for this derivation procedure. In this paper we propose an approach for audit data representation that is geared to simplify the analysis process for the signature engineer. For this purpose audit data and existing relations between audit events are graphically represented. Using a prototype implementation of the approach the strengths of this form of presentation are demonstrated.

Keywords: Computer Security, Intrusion Detection, Misuse Detection, Attack Signatures, Computer Forensic, Data Visualization

Exploring and Modeling the Local Behavior of Personal Machines

Mike Sips (MPI für Informatik - Saarbrücken, DE)

With the increasing availability of personalized web-services running on local machines botnets, trojans, and other kind of malicious code become serious security threads.

These new tools just requires basic programming and scripting skills by the hacker to exploit vulnerabilities of personal services such as IRC, P2P application to infect other machines and launch attacks while an administrator has to invest huge efforts to detect infected machines.

In my talk I will discuss steps toward a integrated framework to support the administrator to detect the point when a local machine shows suspicious behavior. In my research I investigate the idea that knowing the local behavior makes it hard for the attacker to act undetected. I will briefly discuss the current situation, our design decisions and the interactive visual interface.

Joint work of: Sips, Mike; Simon, Sascha; Gerth, John

Interactive Exploration of the Network Behavior of Personal Machines

Mike Sips (MPI für Informatik - Saarbrücken, DE)

Personal machines are often the weakest points within a large network. Although they run an ever-increasing number of network services, these machines are often controlled by users who are unaware of security threats. Thus, a well-informed attacker can, with modest effort, identify and gain control over personal machines. However, system administrators need to know the tools and techniques used for attacks while simultaneously needing to invest huge analytical efforts to detect malicious behavior in the vast volumes of network traffic. In our research project we investigate the idea that an understanding of the regular behavior of personal machines can improve the chance of detecting the point in time when a machine shows malicious behavior. We propose a visual exploration system based on a data abstraction layer and temporal visual representations of the network traffic. The data abstraction layer enables an interactive change in the level of detail of the network traffic while temporal visualizations help system administrators to detect unexpected network traffic. In the next phase of this project, we will conduct experiments to get a good feel about the limits of our system in detecting malicious behavior in real-world scenarios.

Keywords: Visualization, Communication Patterns, Data Abstraction, Personal Machines

Joint work of: Sips, Mike; Simon, Sascha; Gerth, John

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2009/2156>

Visualization of Large Network Structures: Bundled Edges or Node-Link Layouts?

Alexandru C. Telea (University of Groningen, NL)

Visually investigating large network-like structures is a challenging task.

Several approaches have been proposed in the past: node-link diagrams, adjacency matrices, and, more recently, hierarchical edge bundles. In this talk, we present a recent experiment that compares the effectiveness of the classical node-link diagrams with the more recent hierarchical bundled edges. The users involved several computer science practitioners, the data ranged from graphs of several hundreds to several tens of hundreds of nodes, the tasks involved answering a number of structural overview as well as detailed questions involved system dependencies.

Keywords: Visualization of large graphs, user studies, node-link diagrams, bundled edges

Comparison of Node-Link and Hierarchical Edge Bundling Layouts: A User Study

Alexandru C. Telea (University of Groningen, NL)

Visually investigating large network-like structures is a challenging task.

Several approaches have been proposed in the past: node-link diagrams, adjacency matrices, and, more recently, hierarchical edge bundles. We present a recent experiment that compares the effectiveness of the classical node-link diagrams with the more recent hierarchical bundled edges. The users involved several computer science practitioners, the data ranged from graphs of several hundreds to several tens of hundreds of nodes, the tasks involved answering a number of structural overview as well as detailed questions involved system dependencies.

Keywords: Graph visualization, user studies, software visualization, call graphs

Joint work of: Telea, Alexandru; Ersoy, Ozan; Hoogendorp, Hessel; Reniers, Dennie

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2009/2154>

Network Visualization

Jarke J. Van Wijk (TU Eindhoven, NL)

An overview of Network Visualization is given, and the different approaches from graph drawing, information visualization and visual analytics are described.

Keywords: Network visualization, graph visualization, tree visualization