04411 Abstracts Collection
Service Management and Self-Organization in IP-based Networks
— Dagstuhl Seminar —

Matthias Bossardt[1], Georg Carle[2], David Hutchison[3], Hermann de Meer[4], and Bernhard Plattner[5]

[1] ETH Zürich, Switzerland
bossardt@tik.ee.ethz.ch
[2] University of Tübingen, Germany
carle@informatik.uni-tuebingen.de
[3] Lancaster University, UK
dh@comp.lancs.ac.uk
[4] University of Passau, Germany
demeer@fmi.uni-passau.de
[5] ETH Zürich, Switzerland
plattner@tik.ee.ethz.ch

Abstract. From 03.10.04 to 06.10.04, the Dagstuhl Seminar 04411 "Service Management and Self-Organization in IP-based Networks" was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. 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. Service management, network service, self-organization, network management, programmable network, active network, peer-to-peer network, ad-hoc network

04411 Preface – Service Management and Self-Organization in IP-based Networks

Matthias Bossardt (ETH Zürich), Georg Carle (University of Tübingen), David Hutchison (Lancaster University), Hermann de Meer(University of Passau), and Bernhard Plattner (ETH Zürich)

Preface to the online proceedings of Dagstuhl Seminar 04411

Keywords: Service management, self-organization, network management

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/82

Gossip-based self-managing services for large scale dynamic networks

Ozalp Babaoglu (Università di Bologna)

Modern IP networks are dynamic, large-scale and heterogeneous. This implies that they are more unpredictable and difficult to maintain and build upon. Implementation and management of decentralized applications that exploit these networks can be enabled only through a set of special middleware services that shield the application from the scale, dynamism and heterogeneity of the environment.

Among others, these services have to provide communication services (routing, multicasting, etc.) and global information like network size, load distribution, etc. The goal is not to provide abstractions that hide the distributedness of the system, but rather, to hide the unpleasant features of the system, such as dynamism, scale and heterogeneity. Most importantly, these services have to be self-managing: they have to be able to maintain certain properties in the face of extreme dynamism of the network. In this manner, such services can serve as the lowest layer that makes possible building more complex applications, or simply as a plugin to enhance existing systems, for example, GRID environments. Apart from self-management, we require that the services be simple and lightweight, to allow easy implementation and incur low cost.

Our approach to achieving these goals is based on the gossip communication model. Gossip protocols are simple, robust and scalable, besides, they can be applied to implement not only information dissemination, but several other functions, as we will show. So far, we have designed gossip-based protocols for maintaining random overlays, which define group membership. Based on this random overlay, we have designed gossip-based protocols to calculate aggregate values such as maxima, average, sum, variance, etc. We have also developed protocols to build several structured overlays in this framework, including superpeer, torus, ring, binary tree, etc. These protocols build on the random overlay and also on aggregate values. The gossip-based model is well suited to dynamic and large networks. Our protocols are extremely simple to implement while being robust and adaptive without adding any extra components or control loops.

Our approach also support composition at a local level. At each node in the network, the same services are available: for example, data aggregation uses the random overlay (peer sampling service) and superpeer topology construction applies aggregate values, such as maximal and average capacity. In fact, protocols that implement the different services are heavily interconnected and form a modular system within this lighweight self-managing service layer.

While this presentation focuses on the self-managing systems services, it is clear that other application-level services can also be built at higher layers. These services can be proactive, like load balancing, that can make use of the target (average) load and overlays for optimization of load transfer, or reactive, like broadcasting or search, that can be performed on top of an appropriate overlay

network (eg spanning tree or superpeer network), maintained by the lighweight self-managing systems services.

Keywords: Autonomic computing, self management, middleware, aggregation, topology management.

Joint work of: Babaoglu, Ozalp; Jelasity, Mark; Montresor, Alberto

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/81

Service Specific Overlay Setup and Maintenance using Pattern-based Management

Marcus Brunner (NEC Europe - Heidelberg)

Service specific overlays are overlay networks build and setup for a single service. Since different services need different types of overlays the setup process is pretty different. Current approaches normal allow to send in a overlay specification into the overlay provisioning system (mainly centralized). In our approach, we use pattern-based management in order to run the process completely distributed but still having a basic communication pattern to be reused for other overlays with different requirement. In this talk, we present the concept and a first implementation on the Simpson pattern simulator of KTH.

Keywords: Overlay networks, management/provisioning of overlays

Joint work of: Asmare, Eskindir; Schmid, Stefan; Brunner, Marcus

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/83

Toward a More General Network Layer

Ken Calvert (University of Kentucky)

We are developing an integrated network layer capable of supporting unicast, multicast, and publish-subscribe services on any given topology using a single mechanism. Such a service would subsume many of the special-purpose approaches (e.g. overlays) currently proposed and used, and would support novel applications and mobility in a more unified and direct way. A major challenge is to design enrollment, routing,and forwarding algorithms in such a way that the system can scale up to support networks much larger than the current Internet, and to achieve this with overall performance, administrative, and operational overhead comparable to existing approaches.

Our approach is based on the use of predicates carried in packets to identify packet destinations. That is, the network defines a set of predicates over end systems; each packet carries a predicate from this set, suitably encoded, and the job of the network is to deliver the packet (on a best-effort basis) to all end

systems that satisfy its destination predicate, and to as few others as possible. The current Internet is a restricted instance of this model, in which (unicast) addresses correspond to predicates that are satisfied by at most one node.

A key consideration in designing such a system is the relationship between the set of predicates and the topology. Scalability requires the ability to characterize the nodes in a particular region of the network with an efficiently-representable predicate. Supporting mobility or publish-subscribe services, however, may make this difficult. A quantitative notion of "locality" is needed to characterize the relevant tradeoffs.

This talk will discuss a preliminary routing and forwarding design, focusing on the ability of the network to self-organize without pre-assigned addresses on nodes.

Keywords: Routing, forwarding

Joint work of: Calvert, Ken; Poutievski, Leonid; Griffioen, Jim

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/84

Attack Detection using Cooperating Autonomous Detection Systems

Georg Carle (Universität Tübingen)

Today's communication networks are threatened by an increasing number of intrusion attempts, worms, and denial of service (DoS) attacks. Apart from general measures for attack prevention, the possibility to detect ongoing attacks in order to take appropriate countermeasures constitutes an important asset for network security.

A novel approach for attack detection is presented that is based on cooperating autonomous detection systems. The architecture consists of individual detection system that are able to identify ongoing attacks autonomously, and that use cooperation with other detection systems located in other parts of the network for improving the detection performance. Detection of attacks is performed by combining anomaly and knowledge-based detection mechanisms.

Keywords: Autonomous Systems; Attack Detection

Joint work of: Carle, Georg; Dressler, Falko; Münz, Gerhard

Some lessons from an experience with active video flow regulation

Ken Chen (University of Paris-Nord)

People are paying more and more attention to network infrastructures which are capable of dynamic code deployment and reconfiguration, in order to deal with the increase of network complexity both on scale and on heterogeneity. The concept of "active network" has been one of the pioneer ideas. As a starting point, we present an experience we got through the design and implementation of an active network technology based mechanism for video flow regulation. This mechanism makes use of several typical active networking features to perform real-time video flows analysis and provide consequently responsive feedback control to video codec. The main goal here is to adapt quickly the video stream bitrate to the current available bandwidth. From the end-user's view point, the effect of adaptation is to spread the bitrate reduction (relatively) uniformly to all the stream, avoiding in this way abrupt image deterioration (mosaics) due to packet loss. Tests show visible improvements obtained by our mechanism vs the classical RTCP-based control scheme. This work has been jointly done with Rim Hammi. We then discuss some extensions of our mechanism, which is in fact a generic network observer and decision maker.

A more fundamental issue that we identified from this experience is related to the setting of the criteria for code acceptation. This is in fact a rather generic problem, and one can address it in various way. For instance, one can decide to accept a code based on some authentication rule. We are particularly interested by the issue of resources consummation. Indeed, as an example, the network observer module we designed can be configured to get a more or less fine time granularity, and consequently consume more or less CPU. So, one question is how to prevent abusive (either erroneous or malicious) resource consummation. There is few tentative which try to deal with the resource requirement (bandwidth, CPU, memory, etc.) of a code.

The problem is rather complex and hard. It should at least include the monitoring of resource consummation. It requires also a kind of virtual resource model for coding purpose. This issue is, in our opinion, very important. Indeed, we do need a control framework to guarantee not only the correct functionality but also the adequate resource consummation of various codes, in order to be able to deal with future's flexible and/or autonomic networks in a secure and trustable way. Our current research effort on this issue is carried on within the french RNRT/Amarillo project.

Keywords: Active network, video streaming, resource control, responsive control

Joint work of: Chen, Ken; Hammi, Rim

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/85

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/85

Bio-inspired mechanisms for efficient and adaptive network security mechanisms

Falko Dressler (Universität Erlangen)

In recent years, many efforts have been made in developing algorithms and methodologies for building efficient network security mechanisms. The primary requirements are efficiency, adaptability, and scalability. Network security mechanisms are composed of several components. First, high-performance network monitoring entities are required allowing the analysis of transmitted data even in high-speed backbone networks. Secondly, algorithms to detect various kinds of threats have to be developed. Based on the monitored data, statistical anomaly detection methods and policy-based filters can be employed. Finally, the control loop must be closed by involving firewall devices against ongoing attacks.

Organic computing is attempting to build high-scalable architectures, which are self-organizing, self-maintaining, and self-healing. We try to study the processes in computer networks using mechanisms known from molecular biology as the key paradigm. This novel approach shows many similarities between computer networking and cellular mechanisms. Based on the knowledge about cellular metabolism, new concepts for the behavior patterns of routers, monitor systems, and firewalls can be deduced and the efficiency of individual sub-systems can be increased.

This work focuses on the area of network security as one research area with high demand for high-scalable mechanisms providing the needed functionality. We see the proposed mechanism as a generic approach for self-organizing, i.e. self-configuring, self-managing, self-healing, and adaptive solutions in computer networking.

Keywords: Bio-inspired networking, self-organization, network security, organic computing

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/87

Service Provisioning Framework for Self-Organized Networks

Karoly Farkas (ETH Zürich)

Mobile ad hoc networking, as a typical example of self-organized networks, is an emerging and promising communication paradigm. Not only the variety of devices but also the diversity of services is continuously increasing. Such services must be provisioned in a flexible and distributed way without central infrastructure.

Thus, service deployment and management for such mobile devices are extremely difficult since a provisioning framework must cope with the high level of heterogeneity, degree of mobility, and take limited device resources into account.

In this talk, we introduce SIRAMON, a generic, decentralized service provisioning framework for self-organized networks. SIRAMON integrates the required functions to deal with the full life- cycle of services. SIRAMON offers sufficient capabilities to specify, deploy, instantiate and manage not only trivial but also complex services like mobile ad hoc group applications.

Keywords: Service Provisioning, SIRAMON, Ad Hoc Networks, Self-Organization

Joint work of: Farkas, Karoly; Ruf, Lukas; Plattner, Bernhard;

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/88

Deployment Mechanisms for Wide-Area Application Service Overlay Networks

Michael Fry (University of Sydney)

As the Internet evolves the distributed computing model is providing greater opportunities for the provision of new services. Active service approaches, such as funnelWeb, enable dynamic deployment and mangement of new services. However structured management to optimally connect distributed nodes is required. This presentation outlines a set of mechanisms for the discovery of nodes participating in service overlays, and the deployment of application-level services to join specific overlays. The mechanisms are presented within the framework of the funnelWeb infrastructure. Experimental deployment on PlanetLab has demonstrated scalability, robustness and resiliency. These mechanisms are described and illustrated with two examples of application-specific implementations: a E.164 structured overlay, and one based on Global Network Positioning (GNP). The latter will support a variety of services that rely on proximity to nearby nodes on the Internet. We further explore a generic overlay construction mechanism which permits service nodes to mesh optimally with other nodes according to a service-specific metric.

Joint work of: Fry, Michael; MacLarty, Glenn

Peer-to-peer techniques for self-organization in networks

Thomas Fuhrmann (Universität Karlsruhe)

Typical peer-to-peer networks employ self-organization techniques to create special purpose overlay networks, e.g. the overlay of all nodes interested in sharing MP3 files. This talk discusses how the learnings from peer-to-peer networks could be carried over to other self-organization tasks in networking.

Project Venezia-Gondola (A Framework for P-Commerce)

Raymond Gao (P2P Journal - Garland)

A novel project named Venezia-Gondola (Project V-G) was presented, which describes an application platform that enables the activities of Peer-to-Peer commerce (P-Commerce). A new pattern called the Inverted Model-View-Controller (IMVC) pattern was claimed that is suitable for P-Commerce. The author also explains the principles of the Project V-G and possible architecture for future development.

Keywords: Peer-to-Peer Commerce, Ad-Hoc Transaction, JXTA, XML, Design Patterns

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/89

Charging Service Compositions in a Service-Oriented Peer-to-Peer Network

Jan Gerke (ETH Zürich)

Today, peer-to-peer (P2P) networks, e.g., filesharing networks like Gnutella, are specialised towards specific purposes. This shortcoming is addressed by introducing a new middleware for P2P networks. The middleware supports the deployment and use of services inside a P2P network.

The middleware can be adapted through plug-ins. It includes service negotiation mechanisms which support the creation of legally enforceable service level agreements (SLAs) by using strong identities. This allows to reliably compose services into new value-added services.

The business model of service composers has been investigated. The question of how to charge for such service compositions is addressed by introducing an event-based charging model. The events are to be applied in a service consumer's utility function as well as service composers' and service providers' tariff functions. Finally, the relation of these functions determines the profits of a service composer.

Keywords: Peer-to-peer middleware service negotiation service level agreements service composition service charging

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/90

Whither Service Management - AlbatrOSS and Beyond

Jane Hall (FhG FOKUS - Berlin)

Service management is facing new challenges given the evolution in new generation technologies, the service paradigm, and business models for a multi-service and multi-provider environment. A 3G OSS architecture for a personalised mobile services environment developed in the AlbatrOSS project is examined in order to illustrate these issues and indicate directions forward. A continuation of the work is intending to investigate management aspects of service modelling and service composition.

Efficiency and Quality of Service of IP Service Provisioning

Oliver Heckmann (TU Darmstadt)

The work presented consists to a large extend of the dissertation.

A high efficiency (i.e., the ratio between the transported traffic and the costs for transporting the traffic) and a well-defined quality of service are two very important goals for IP networks. From this follows the question: How can an Internet service provider optimise the efficiency and quality of service of its network rsp. of service provisioning?

This work shows that efficiency and quality of service strongly depend on the network architecture, traffic engineering, network engineering, and the interconnections. Therefore, a system-oriented approach is employed. With this approach, all previously mentioned areas are analysed and optimised, considering their mutual influences.

In the context of network architectures, the focus of the research lies on methods for providing quality of service. A bandwidth broker for Differentiated Services networks is developed. It offers strong quality of service guarantees and a high efficiency by overbooking.

Besides the network architecture, the interconnection structure (i.e., the connections to other networks) also influences the efficiency and the quality of service of a network. In this work, different approaches to optimise the efficiency, reliability, and quality of service of the interconnection structure are developed and evaluated. The work shows that with these approaches, significant cost savings and quality of service improvements can be realised.

How the efficiency respectively the quality of service obtained with the network architecture and interconnections can be further improved by traffic engineering is also analysed within this work. Weaknesses of existing approaches are identified and corrected.

Due to the steadily increasing traffic volumes, capacity expansion is the most important network engineering task of a provider. Therefore, capacity expansion strategies are elaborated and evaluated in this work. They also consider the influence of different network architectures and traffic engineering.

Keywords: QoS, efficiency, IP service provider, ISP, optimisation, peer-to-peer

Enriched Classification and Dynamic Tunneling for Supporting Mobile IP Services

Gísli Hjálmtýsson (Reykjavik University)

We make the case that mechanisms for enriched classification and dynamic tunneling with a common elementary protocol is both necessary and sufficient for facilitating mobile IP services in highly environments. Moreover, we claim that device initiated handoffs simplify the network infrastructure and are in general more suitable in highly mobile and heterogeneous environments. We discuss our preliminary work to validate these claims based both on simulation models and experimental prototyping.

Keywords: Protocol design, Tunneling, Classification

Joint work of: Hjálmtýsson, Gísli and Brynjúlfsson, Björn

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/91

Service creation, composition, and deployment in mobile and pervasive environments

Mohan Kumar (University of Texas at Arlington)

The gap between computing and communication devices is decreasing rapidly Ű in the near future every device will be required to perform both types of services. Network processors, active networks and ad hoc networks subscribe to this trend. Proliferation of wireless communications would require all devices to serve as communication routers to switch packets among communicating and computing devices. The number and variety of devices that are being used in mobile wireless environments and various application domains are increasing at an enormous rate. However, many of these devices such as sensors, embedded processors, personal servers, cell phones and PDAs have limited resources in terms of CPU, memory, communication bandwidth, and battery life. Therefore, it is extremely important as well as challenging to conserve device resources, yet provide software services in a timely fashion for a variety of application scenarios. Service creation, composition and deployment will be integral to computing environment environments. Service composition, deployment and provisioning will demand robust, but adaptable communication networks. Active networking technology is the ideal framework for supporting network adaptability in dynamic heterogeneous environments. In this seminar, we present ongoing research work on services in pervasive computing environments and the role of active networks.

Keywords: Network services Reource-limited devices Pervasive Computing Active Networks

Self-Organization in Peer-to-Peer Systems

Hermann de Meer (Universität Passau)

Peer-to-Peer Systems are about community-based cooperations. The peers share responsibilities and benefits by cooperating in a distributed and decentralized environment. To carry out tasks sensibly, however, a more or less rigid order is required for efficiency and reliability reasons. This order can be partially imposed from the outside, for example within so-called "structed" Peer-to-Peer systems. A common approach here is the use of Distributed Hash Tables. Alternatively, Peer-to-Peer systems can be "unstructured" in the sense that an useful order emerges from own internal processes. Unstructured and structured Peer-to-Peer systems rely both on a more or less decentralized overlay management. Self-organization, therefore, is a key to the success of Peer-to-Peer systems in various forms. This presentation gives an overview of the role of self-organization in Peer-to-Peer systems.

Keywords: Self-organization, peer-to-peer

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/86

Composite Protocols and Networking Services

Gary Minden (The University of Kansas - Lawrence)

Active Networking is concerned with the rapid definition and deployment of innovative, but reliable and robust, networking services. Towards this end we have developed a composite protocol and networking services architecture that encourages re-use of protocol functions, is well defined, and facilitates automatic checking of interfaces and protocol component properties. The architecture has been used to implement common Internet protocols and services. We will report on this work at the workshop.

Keywords: Active Networking; Composite Protocols; Composite Network Services

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/92

UMM: A dynamically adaptive, unstructured multicast overlay

Matei R. Ripeanu (University of Chicago)

The simplicity of multicast as a communication primitive belies its broad utility as a building block for distributed applications. Nevertheless, creating and maintaining multicast structures can be challenging, particularly when networks are transient and/or dynamic. We introduce a new unstructured multi-source multicast (UMM) overlay approach that we argue is less complex than, but as efficient as, current state-of-the-art solutions based either on structured overlays or on running full routing protocols at the overlay level. UMM builds a base overlay independently from the routing mechanisms employed to route messages. On top of this base overlay, it selects distribution trees for each multicast source by first flooding the base overlay and then using the implicit information contained in duplicated messages to select and filter out redundant tunnels. Simple heuristics are used to maintain and evolve both the base overlay and the multicast distribution trees in response to changes in the set of overlay participants or in underlying network conditions. We experiment on a 65-node PlanetLab deployment and on ModelNet emulated distributed platforms to quantify the overheads associated with UMM operation and to explore its performance and adaptability to changes in the underlying network conditions.

Keywords: Multicast overlay, self-organization

Joint work of: Ripeanu, Matei; Foster, Ian ; Iamnitchi, Adriana; Rogers, Anne

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/93

A flexible router platform for next generation network services

Lukas Ruf (ETH Zürich)

Autonomous services need a flexible router platform that provides the mechanisms to install, modify and remove services at run-time of the node without interfering with others. Instantiated services must have the ability to re-configure and to exchange service functionality themselves. Envisioned router platforms must be able to run multiple services in parallel and are required to scale with the number of network-interfaces while they need to provide a straightforward to use service programming interface.

In this talk, we present the PromethOS NP router platform together with a service architecture to counteract distributed denial of service attacks in an autonomous way. PromethOS NP manages and controls a processor-hierarchy

composed of host processors and network processors embedded in network interface cards. It provides a dynamically code-extensible router platform of which all processor tiers are at run-time programmable following a unified component programming model.

The service architecture illustrates the capabilities of the router platform and its applicability to autonomous network services.

Keywords: Extensible Router Platform, Service Architecture, Network Processor, Host Processor, Processor Hierarchy

Joint work of: Ruf, Lukas; Wagner, Arno; Farkas, Karoly and Plattner, Bernhard

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/94

DotSlash - Creating Content Distribution Networks on Demand

Henning Schulzrinne (Columbia University)

Traditional content distribution networks, such as Akamai, are well-suited for static web services that routinely experience large traffic volumes. They are unsuited for active content, i.e., content generated by scripts from databases, and web sites that are unlikely to receive significant number of requests. However, a few such sites will invariably experience their "fifteen minutes of fame", typically by being mentioned on a high-volume news site such as SlashDot or CNN. Such flashcrowds or "slashdot effect" will routinely cause single-server websites to collapse.

We have designed and prototyped an autonomic web replication system, called DotSlash, that drafts rescue servers fully automatically, without user intervention. The system discovers suitable rescue servers via wide-area service location, either among peer servers or from a dedicated pool of rescue servers, allocates them for temporary use and redirects requests to them. The system is completely transparent to clients and does not require URL rewriting or other client modifications.

We have designed two versions. The first, an Apache extension, deals only with static content, e.g., HTML pages or media objects. The second version can also replicate and execute scripts remotely. We have prototyped the system for the common LAMP (Linux, Apache, MySQL and PHP) configuration and shown that a common benchmark for bulletin boards can be replicated without code changes, yielding capacity increases bounded only by the database server. Since many such systems, including most blogs, are bottlenecked by the web server, our system can significantly increase capacity and works even for extremely rapid load increases.

We are currently investigating how such systems can be further extended by increasing the database capacity of read-mostly systems with loose consistency constraints.

Keywords: CDN; web replication; flash crowd; Slashdot effect; autonomic system; PHP; LAMP; active content

Joint work of: Zhao, Weibin; Schulzrinne, Henning

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/95

From Active Networks to Cognitive Networks

Manolis Sifalakis (Lancaster University)

Future networks need to be autonomic self-managed and provide resilient servicing, even when the hardware fails. To achieve this goal, two fundamental requirements need to be satisfied: (i) the service management and provisioning must be independent and decoupled of the infrastructure management, and (ii) a certain degree of cognitive behaviour needs to be achieved at the service management level. In achieving the first goal, which in turn will enable the pursuing of the second goal, active and programmable networks will play an important role. A problem though arises when we try to build and use actual active networks, as most research so far has focused at the node level and has left us with a unbridged diversity of platforms and execution environments, which are largely uninteroperable with each other. We introduce a toolkit that provides a set of mechanisms aiming to bridge this diversity and provide a set of functionalities and abstractions for uniform installation and deployment of services over active and programmable networks.

Keywords: Active networks, self organisation, service mobility, dynamic service deployment

Joint work of: Sifalakis, Manolis Hutchison, David

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/96

Peer-to-Peer vs. the Internet: A Discussion on the Proper and Practical Location of Functionality

James P. G. Sterbenz (Univesity of Massachusetts - Amherst & Lancaster University)

Peer-to-peer information sharing has become one of the dominant Internet applications, measured not only in the number of users, but also in the network bandwidth consumed. Thus, it is reasonable to examine the location of support functionality such as self-organisation, resource discovery, multipoint-to-multipoint group communication, forwarding, and routing, to provide the needed service to applications while optimising resource usage in the network.

This position paper is intended to stimulate discussion in two related areas: First, where *should* functionality to support peer-to-peer applications be located:

in the network, or as an application overlay among end systems. Second, where *can* functionality be located, given the practical constraints of the modern Internet including closed systems and middleboxes, as well as administrative, legal, and social issues. We will discuss the performance implications of these decisions, including whether low latency bounds for delay sensitive peer-to-peer applications (such as distributed network computing) can ever be achieved in this environment.

Keywords: Network architecture, peer-to-peer, client/server, end-to-end arguments, protocol layering, policy, ussle

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/115

Self-Healing Protocol Implementations

Christian Tschudin (Universität Basel)

Current studies on self-configuring and adaptive networks aim at developing specific and fixed protocols which are able to optimize their configuration in a variable network environment. In this talk we study the problem where the protocols need to cope with a defective execution, including the lossy execution or the injection of foreign code. One guiding question will be the creation of robust execution circuits which can distribute over a network and which continue their service despite parts of the implementation being knocked out. The ultimate goal is to enable protocol implementations to detect by themselves that they are malfunctioning and to let them correct their own operation mode and code base.

As a show case, we present a protocol implementation which is robust against deletion (knock-out) of any single instruction, regardless whether this deletion affects the core protocol functionality or the resilience logic. The technique used in this first of its kind example is the self-modification of the running program, which can be naturally situated in an active networking context. Ultimately, a self-correcting protocol implementation has to constantly rewrite itself according to the (self-)observed performance.

In this talk we will also point to related fields like self-correcting software, fault tolerant quantum computing and self-healing properties of biological systems.

This is joint work with Lidia Yamamoto, Hitachi Europe.

Keywords: Resilient protocol implementations, autonomic communication systems, active networking, bio-inspired networking

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/98

Simulative Performance Evaluation of a Mobile Peer-to-Peer File-Sharing System

Kurt Tutschku (Universität Würzburg)

Peer-to-Peer (P2P) file-sharing has become the killer application in the wired Internet and might also be highly attractive for mobile networks. In particular since UMTS operators are searching for new applications which do both: a) exploit the potential of the UMTS technology and b) motivate the user to adopt the new technology.

In this work we are investigating the performance of an eDonkey-based mobile P2P file-sharing system by means of time-dynamic simulation. Mobile networks differ from wireline networks by the limited capacity of the radio link and the mobility of the users. P2P networks, in contrast, are overlays which consider the transport network in an abstract way. In a mobile environment, the question arises, whether the abstraction can be maintained and what will be the performance impact if there is any. We will show in detail how the mobile access technology (GPRS or UMTS), the churn behavior of mobile users, the file size of mobile specific content, and special infrastructure entities, such as a cache peer, influences the performance of the suggested mobile P2P file-sharing service.

Keywords: Peer-to-peer, UMTS, performance evaluation, file-sharing

Joint work of: Tutschku, Kurt; Hoßfeld, Tobias; de Meer, Hermann; Oberender, Jens; Andersen, Frank-Uwe

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/99

Policy Management of Pervasive Environments: Lessons for Network Architectures

Ian Wakeman (University of Sussex)

In this talk, I outline our approach to policy management of pervasive computing contexts. We believe that users will best be able to cope with the complexities of policy configuration if they can describe interfaces in their own natural language. To this end, we are building a natural language interface which uses description logic as the policy formalism. I descibe the middleware that results from using description logic, and outline some of the advantages that accrue from using a high level logic in policy analysis and execution. I conclude with some outrageous claims about what the network should look like if computing contexts become widespread.

Keywords: Policy management, natural language processing

Joint work of: Keller, Bill; Owen, Tim; Wakeman, Ian; Weeds, Julie; Weir, David

Bridging today's Internet Heterogeneity with a Content-oriented Approach

Klaus Wehrle (Universität Tübingen)

Despite the success of Internet technology and protocols, their deficiencies become apparent with growing and emerging applications like 3G mobile devices or pervasive environments. In these fields, there is a growing demand for additional support of mobility, multi- or anycast communication, or service composition.

Keywords: Peer-to-Peer, Overlay, Communication Form, Bridging Heterogeneity

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2005/100

Epidemic Dissemination in Ad Hoc Networks

Lidia Yamamoto (Hitachi Europe - Valbonne)

Peer-to-Peer (P2P) and ad hoc networks have many points in common: both represent a decentralized self-organizing network structure. However few existing P2P algorithms are specifically designed to operate efficiently over ad hoc networks. And few ad hoc networks are designed to benefit from P2P infrastructures.

We have worked on an epidemic dissemination protocol to maintain soft-state in a decentralized, peer-to-peer fashion, in ad hoc networks. This protocol is an enhancement of Passive Distributed Indexing (PDI) method proposed by Lindemann and Waldhorst. PDI is a method for distributing information in a P2P structure which is particularly suited to ad hoc networks, and does not involve an overlay topology. It makes use of broadcast messages to spread information via passive epidemic dissemination.

We have enhanced PDI in order to reduce the number of broadcast messages when the search for an item may span several hops. Three enhancements are proposed: 1) Lazy query propagation to delay the propagation of query messages such that local responses can inhibit unnecessary search. 2) Quench waves to stop an already initiated query propagation when still possible. A decision algorithm determines whether to start a quench wave or not based solely on local information. 3) The use of Multi-Point Relay (MPR) or similar protocol and algorithm, to reduce redundant broadcast messages.

This talk will present the current state of this research, and discuss several open aspects with the purpose of stimulating debate.

The talk will also include an overview of related work such as epidemic models from biology, other epidemic protocols for P2P overlays and MANETs, including gossip (active) and promiscuous (passive) dissemination modes.

Such protocols could be used for many different purposes, roughly any task requiring distributed soft-state maintenance in the ad hoc network, including

DNS and identifier mappings, network monitoring and configuration, and so on. During the talk we will also exploit the possibility of using the protocol to disseminate service information for on-demand service deployment, and further, to assist in self-composing protocol structures.

Keywords: Ad hoc networks, epidemic dissemination, peer-to-peer

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2005/101