Key Research Challenges in Network Management

Aiko Pras, Jürgen Schönwälder, Mark Burgess, Olivier Festor, Gregorio Martínez Pérez, Rolf Stadler, and Burkhard Stiller

ABSTRACT

Although network management has always played a key role for industry, it only recently received a similar level of attention from many research communities, accelerated by funding opportunities from new initiatives, including the FP7 Program in Europe and GENI/FIND in the United States. Work is ongoing to assess the state of the art and identify the challenges for future research in the field, and this article contributes to this discussion. It presents major findings from a two-day workshop organized jointly by the IRTF/NMRG and the EMANICS Network of Excellence, at which researchers, operators, vendors, and technology developers discussed the research directions to be pursued over the next five years. The workshop identified several topic areas, including management architectures, distributed real-time monitoring, data analysis and visualization, ontologies, economic aspects of management, uncertainty and probabilistic approaches, as well as understanding the behavior of managed systems.

INTRODUCTION

Many researchers come together on various occasions to discuss promising directions for future research. Ideas from such discussions are important to define new and possibly joint research projects, to direct funding organizations such as the European Union (EU) and the National Science Foundation (NSF), and to guide Ph.D. students. In the general area of networking, the NSF has sponsored a number of workshops to identify fundamental research challenges to networking [1]. In the area of network management, the Internet Architecture Board (IAB) has organized a workshop to guide the Internet Engineering Task Force (IETF) in their work on standardizing network management protocols [2]. The IAB workshop has been quite successful, and it paved the way for new protocols such as network configuration (NETCONF). However, the focus of that workshop was on standardization and not on fundamental research.

In October 2006, the Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF) organized a workshop to define challenges for future research in the area of network management. The workshop was organized together with the European Network of Excellence (NoE) for the Management of Internet Technologies and Complex Services (EMANICS), which is supported by the European Commission and investigates how to manage the future Internet, including the services running on top of it, in a scalable, economic, and automated way. The workshop was held on the premises of SURFnet in Utrecht, the Netherlands and attended by twenty representatives from academia, vendors, and operators. The number of participants was deliberately kept low to foster discussion. Prior to the workshop, an open call was sent via the NMRG mailing list for position statements; the selection of participants was based on these position statements. Because it is important for researchers in the area of network management to interact with operators and vendors, the workshop organizers ensured that there would be a good balance between these groups and that participants would come from all parts of the world. The precise list of participants is contained in the meeting minutes and can be downloaded from the NMRG Web site. The first day of the workshop was used to present and discuss the various position statements; at the beginning of the second day, the group divided into parallel vendor, operator, and researcher sessions. At the end of the workshop, the three groups came together again to discuss their findings and to draft conclusions.

The goal of this article is to summarize some of the main discussions at the workshop. It should be noted that the article presents only a selection of the discussions; more details can be found in the meeting minutes. Also, it is important to note that this article reflects the research interests of the workshop attendees only. Although the call for position statements was open to everyone, some important researchers were unable to attend. As a consequence, some topics may not have been discussed in as much depth as would otherwise have been the case. An example of such topics

0163-6804/07/$20.00 © 2007 IEEE

is the topic of management policies.

The remainder of this article is structured according to the following topics: management architectures, distributed monitoring, data analysis and visualization, ontologies, economic aspects of management, uncertainty and probabilistic approaches, and finally, the behavior of managed systems.

ARCHITECTURES

In the past, many researchers worked on the definition of network management architectures. As a result, we now have a good understanding of the manager agent model and several forms of distributed management (Fig. 1). The IETF, for example, has been working on three different kinds of distributed management standards: a management information base (MIB)-based approach (expression, event, and notification MIB), a script based approach (script and schedule MIB), and a remote operations based approach. Also, organizations such as the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T), as part of their telecommunications management network (TMN) series of recommendations, have defined functional, physical, information, and logical layered architectures. These architectures have in common that they rely as a basis on a client-server model. Although such models are still useful for many traditional management tasks, they often fall short for managing emerging peer-to-peer (P2P) and ad-hoc networks.

In P2P systems, the scale, dynamics, and differences in business models may be such that it becomes quite difficult for centralized managers to perform traditional management tasks, such as quality of service (QoS) monitoring and fault handling. Therefore, P2P systems require cooperative management capabilities, for example, to collect usage statistics, repair faults, pass network address translations (NATs), and find users. Also, ad-hoc networks might not be manageable from a centralized system. Still, such networks perform tasks that might be considered as cooperative management tasks; examples are the decision of when to join two networks, when and how to split a network in case of overload, how to decide which devices become responsible for external connectivity, and how to deal with security attacks.

Cooperative management often will be performed in an automated way. To stress the reduced role of the human manager, terms like implicit, autonomic, or self-management sometimes are used to denote such automation. A problem with automated management mechanisms, however, is that multiple control loops may be created, which work well in isolation but may interfere with each other in exceptional cases and put the stability of the system in danger. Regaining control in highly dynamic environments, such as P2P and ad-hoc networks, may be far from trivial, and therefore further research is required to investigate stable cooperative management models for these types of (autonomic) networks.

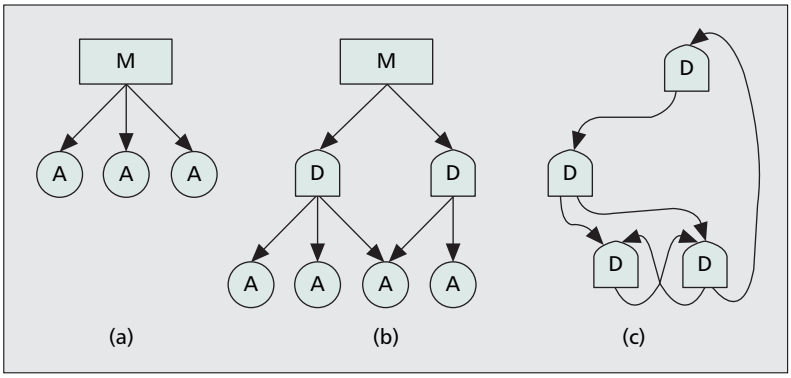

Figure 1. *Different forms of management: a) centralized; b) distributed; c) cooperative.*

DISTRIBUTED MONITORING

Effective monitoring provides state information with the required accuracy to the right places in the network, at minimum cost. For many management functions, including quality assurance, proactive fault, and security-management, monitoring in real-time is required. Some workshop participants argued that to achieve the goal of scalability and fast reaction times for future networks, the monitoring activity should be provided by the network itself. Therefore, a research effort should be directed at engineering a distributed, self-organizing monitoring layer that resides inside the network and offers end-to-end monitoring primitives to management applications and end systems.

A key function that a distributed monitoring layer must provide *is estimating aggregates of variables in real-time*. Such aggregates may be computed across nodes in a neighborhood, a network domain, or the entire network. Simple examples of aggregates include sums, averages, extreme values, percentiles, and histograms of device counters. For the purpose of quality assurance, it may be required to continuously track the number of voice over IP (VoIP) flows in a network domain or the distribution of traffic composition across all links. Similarly, to achieve a given level of availability, it may be necessary to know, at all times, the percentage of links that operate above 50 percent utilization and to identify the 10 most loaded links.

The challenge is to engineer a set of protocols for such a distributed monitoring layer, which performs the tasks of *distributed polling*, *continuous estimation*, and *threshold detection for network-wide aggregates*. To keep the complexity of the monitoring layer low and to enable efficient, effective, and scalable operation, these protocols must be self-configuring, robust, and tunable at run time.

In large-scale networks, continuous monitoring with maximum achievable accuracy of even a single aggregate can become unfeasible, due to high traffic and processing overhead. In addition, modern routers contain hundreds of counters that are locally available for monitoring, many of which are required in aggregated form to support management control loops. Consequently, when designing monitoring protocols, the engineering trade-offs *must be controllable* at invocation or at event run time.

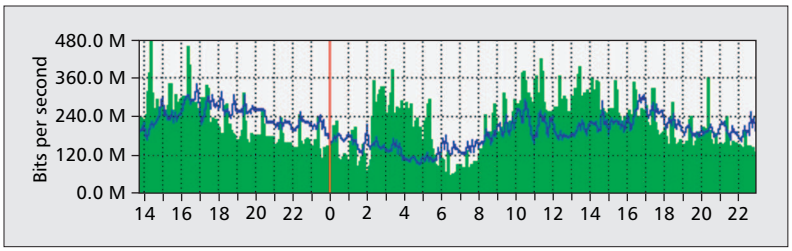

Figure 2. *MRTG graph — example of time series plot.*

Monitoring protocols can be optimized toward providing estimates with low overhead, small delay, high accuracy, or a high degree of robustness. Because jointly optimizing these metrics generally is not possible, the right operating point in the parameter space created by these metrics must be selected. For instance, recent results suggest that allowing for some modest errors, the protocol overhead to estimate an aggregate can be reduced by an order of magnitude in a realistic setting [3]. Taking into account that different management applications have different requirements regarding the quality of the estimates (delay, accuracy, etc.), the operating point must be a control parameter of the protocol. By allowing a management application to change the operating point at run time, monitoring functions can be built that adapt their operation to the required quality of monitoring data, which may change over time.

A known approach to compute aggregates in a distributed fashion involves building a tree-based overlay in the monitoring layer and aggregating state information along that tree, bottom-up from the leaves towards the root. Such trees can be built in a decentralized, self-stabilizing manner that provides the monitoring protocol with robustness properties. Recently, it has been suggested to use gossip protocols, which typically rely on randomized communication to disseminate and process information in a network. Although many believe that gossip-based monitoring is likely to be more robust than tree-based monitoring, the problem of "mass loss" that is intrinsic to a gossip-based solution must be addressed.

Research into efficient state aggregation under constraints has recently been proposed in different contexts, for example, for wireless sensor networks. For the field of network management, the constraints that are specific to a distributed monitoring layer, the rich functionality of network management operations, and the potentially large number of concurrently executing operations within such a monitoring layer make the problem unique and interesting.

Data Analysis and Visualization

Network management systems collect large monitoring and measurement data sets that must be aggregated, filtered, and visualized with the goal of making meaningful information easily accessible to human network operators.

First generation network management systems were well known for their topological network views. They introduced automatic discovery and mapping procedures to make it simple to set up topological maps. Next to topological views, network management systems usually provide time series plots to visualize the evolution of key metrics over time, where the time scale typically varies from days and weeks to months and years (Fig. 2). Most systems today are accessible via Web interfaces (periodically updating Web pages), but experiments also have been made using TV channels to make network status graphs, sometimes also called network weather maps, accessible to a large number of network users.

Although data analysis and visualization is an old network management topic, it seems that available techniques and interfaces for human network operators are not really satisfying for the following reasons:

- Traditional topological views, especially those based on geographical maps, do not scale well with the growing number of networks and network elements. Furthermore, there is a multitude of different layers involved and attempts to visualize topologies on multiple or all layers makes the scalability problem worse.
- Collected measurement data sets and statistics often are visualized in a rather static way. There is typically no or only very limited support to explore data sets (e.g., by applying filters, zooming functions, or correlation functions) in an interactive way.
- Traffic visualizations typically focus on the visualization of high-volume traffic components or flows. Although this is certainly useful for planning and perhaps accounting purposes, there is also a growing need to extract and highlight the unusual traffic and unusual traffic patterns. Especially for security auditing purposes, it is often much more desirable to find and locate small volume but highly unusual traffic streams or patterns.
- Many existing tools are designed for offline analysis and visualization. There is, however, a growing need for online, close to real-time analysis and visualization, to reduce detection and reaction time. With network transmission speeds of a multiple of tens of Gigabits per second, this becomes non-trivial, as data capturing alone becomes challenging. The move towards statistical data capturing methods for high-speed networks also requires the visualization of the accuracy of data sets.

Some basic research on data visualization techniques has been done as part of the Cooperative Association for Internet Data Analysis (CAIDA) project. The CAIDA project has developed various techniques to visualize the autonomous system (AS) interdomain backbone network, some based on geographic maps and others based on more abstract representations. The Walrus graph visualization tool can be used to visualize large directed graphs in three-dimensional space, which according to the CAIDA project page, is effective for graphs up to a few hundred thousand nodes and only a slightly greater number of links. In addition, CAIDA has

performed work on the visualization of flow data sets, the outbreak of worms, and backscatter traffic.

Research on three dimensional visualization techniques also has been done in several smaller and more focused projects. While many of these visualizations look rather fancy, their usability has not been well analyzed, and the creation of such visualizations usually requires elaborate hardware and software tools. So far these technologies are not widely accessible and usable. Given the recent improvements in graphics capabilities of workstations today and new interactive formats, there is a good chance that effective multi-dimensional visualization tools may be developed in the near future. For geographic maps, we envision that recent technology, such as Google Earth, will act as an enabler for the development of new techniques where zooming and on-the-fly data aggregation can be explored and integrated with widely accessible geographic information systems.

ONTOLOGIES

Another research challenge identified at the workshop was ontologies, in particular the role they play in network and service management and the added-value they provide to this field.

One of the base components in network management is the model used to represent the different elements and objects being managed. In our field, current models usually are classified either as data models (DM) or as information models (IM). The main difference between the concepts is the level of abstraction. Data models are closer to the underlying protocols used to transport the management information and the particular implementation in use. In fact, they are intended for implementers.

In contrast, information models work at a conceptual level, and they are intended to be independent of any particular implementation or management protocol. Working at a higher level, information models usually provide more expressiveness to designers. More information about the difference between data models and information models can be found in [4], which describes the main results of a previous IRTF-NMRG meeting.

Both data and information models are now being complemented by semantic models (Fig. 3), where the meaning of concepts used in the network management field and relationships existing between them are made explicit. Additionally, this meaning is defined in a machine-readable format, thus making it accessible to both software management components and humans.

The use of semantic models eases interoperability between different management domains and applications, not at the level of data exchange, which is mostly solved with standard data models that currently exist (e.g., SNMP MIB), but at the level of knowledge sharing. Achieving this objective enables different administrators and/or management software components to clearly understand the definitions and management rules and goals provided by other administrators (possibly using different management platforms). It also enables the understanding of what the administrators and/or management software components really meant when they provided these definitions and management objectives. In [5], the authors describe one specific application scenario where these semantic interoperability ideas are illustrated for router configuration management.

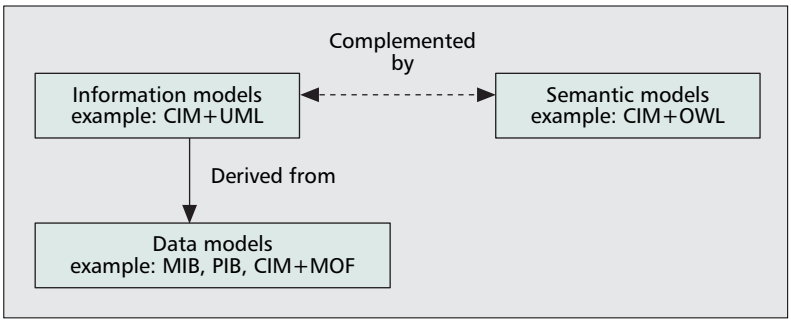

Figure 3. *Data, information, and semantic models.*

To define semantic models in the management field, a formal definition of the knowledge used in this domain is required. Ontologies can model the semantics of managed entities and the relationships existing between them. In fact, ontologies can be defined as a formal, explicit specification of a shared conceptualization [6]. As such, they try to provide a shared and common understanding of a certain application domain, in our case network management, thus facilitating the exchange of information, including well-accepted semantics in the management field.

Ontologies are defined in a machine-readable format, using languages like the ontology Web language (OWL) from the W3C, whose definition can be accessed by either software management components or humans, thus providing an added-value over information models that are mostly visual and usually intended for humans.

The use of this semantic modeling approach also benefits the non-trivial processes related with network management. Examples of these processes are conflict detection and resolution and policy refinement from high-level objectives to low-level configurations. More work is expected in these research areas in the next few years.

Regarding limitations, it should be mentioned that ontologies are still under development in the management field. In fact, the technology is not yet mature, and there is not an ontology that can be considered as a de-facto standard by the international community. However, some research efforts should be noted related to the definition of ontologies based on the common information model (CIM), although these works still must be refined and evaluated in different real networking scenarios. These limitations also provide areas for future research in this field.

ECONOMIC ASPECTS

Network and service management in the traditional view considers mainly technical parameters. This view included parameters that were measured or monitored within the network, at its borders, and sometimes within end-systems of

The importance of uncertainty in management is even growing as a foundation of most emerging networks with high dynamics such as P2P overlays or ad-hoc networks. Today, even management data must be considered with some degree of uncertainty.

users or customers. Those values that were measured were accounted for (e.g., by applying the Remote Authentication Dial-In User Service (RADIUS) protocol or related vendor-specific systems) and maintained in respective databases (e.g., by an implementation of the authentication, authorization, and accounting (AAA) architecture of the IRTF) of operators and their operation support systems (OSS) as well as business support systems (BSS).

An OSS mainly deals with day-to-day operations, fault and error handling, and channel supervision and selection [7]. The optimization of resource usage for medium- and long-term falls more into BSS tasks. These tasks are extended further as soon as business support cycles come into the picture: now the economic viewpoint of network management must be handled explicitly, for example, for service provisioning, service tariffing, defining service mixes for a given or intended customer base, and last but not least, for the cost of network and service management in relation to the economic gains to be achieved for a given network infrastructure. To complete the picture, in addition to technical and economic dimensions, the role of a service level agreement (SLA) and its legal content is highly relevant [8]. SLAs can be enforced, and the commercial user can be provided with a managed service based only on commonly agreed-upon legislation, which today typically follows the territoriality scheme.

Considering network management tasks as a whole, from an economic point of view, they form a certain type of risk management, which follows a predetermined list of objectives, such as avoiding overload situations of network links, minimizing the packet-loss rate in the whole network, or ensuring the availability of network services for premium customers. Thus, all mechanisms that are put in force to achieve such objectives can be considered as an insurance, which must be paid in advance at the time of putting a network into operation and which must be paid during the operation of the network for collecting and relating the right data. Thus, the key economic dimension to this approach is to determine the cost and potential trade-offs between:

- Applying the best suited technical and economic management mechanisms to a network and its commercial users.
- Not managing such a network infrastructure at all.

For a given network infrastructure, an applicable cost model is the key to determine a reliable answer. This model can form the basis for which management costs can be justified, which includes customers as well as vendors of network equipment. The application of traffic management schemes, compared to a purely over-provisioned network, especially give rise to discussions of whether one or the other is more efficient. Therefore, the efficiency in measurements, a possibility to cover service differentiation, and a valuation of the traffic mix determine economic incentives, which will ease the selection of best-suited mechanisms to run network management functionality within a single network and in the multi-domain case. However, it is important to note that all operators must improve their traffic flow and services handling in technical and economic terms.

Based on these brief explanations on the importance of economic and legal dimensions in network and service management, for the commercial user and the services utilized, key issues must be investigated in much more detail for improved management mechanisms.

Although network technology and the dedicated end-to-end service mix offered must provide incentives for operators to collaborate in management tasks or at least interoperate based on standards — typically to avoid negative externalities — these incentives must be made explicit in a fully decentralized management approach. Additionally, modeling the cost of management mechanisms in place today and expected for tomorrow becomes mandatory to enable answering emerging questions such as:

- How many management functions will be integrated into a network?
- Which granularity of measurement and monitoring data must be archived and related to enforce a certain economic objective?
- Which tariff models will be applicable and efficient for multi-domain end-to-end services?

Thus, the core of economic management will enable operators in the future to ensure a complete, coherent, and vertical service offering under integrated technical and economic optimizations.

Uncertainty and Probabilistic Approaches

Probabilistic approaches are used very successfully in almost all scientific disciplines. Their use enables the design of rich models of the studied systems that would be otherwise almost impossible to model. Several recent management approaches exploit this potential. The most well-known are the monitoring approaches using sampling techniques to collect useful data to establish the current state of a system. Probabilistic techniques drastically reduce the cost of management (i.e., the cost of physical probes, collectors, and high performance data storage) and enable the calculation of the accuracy of the measured data. Sampling techniques have proven useful for IP packet-level monitoring. More recently, they also were successfully applied directly to distributed monitoring functions to measure, for example, a near real-time estimate of VoIP flows in a large network [3].

As in many other disciplines, uncertainty is part of daily life in network management. Although uncertainty can be seen as pain that we must deal with, it also can be seen as an enabler, fostering the design of algorithms capable of computing, predicting, and even influencing the managed systems behavior.

The importance of uncertainty in management is even growing as a foundation of most emerging networks with high dynamics such as P2P overlays or ad-hoc networks. Today, even management data (i.e., data collected regarding the state of the network and configuration orders) must be considered with some degree of

uncertainty (e.g., devices can cheat about their state or behavior; received data may be outdated or lost).

The necessity of dealing with uncertainty in the management plane became obvious when management was considered for emerging networks, especially ad hoc networks. To cope with dynamics and scale, distributed management algorithms, using uncertainty as a predicate, were successfully designed and evaluated for selective distributed monitoring, fault-management [9], as well as for configuration management in similar networks.

Several attempts have been made to deploy apparently deterministic methods, with some success. However, algorithmic determinism does not guarantee behavioral determinism. The highly dynamic nature, complexity, and size of network infrastructures today, as well as the strong and often unknown dependencies among elements, events, and management actions do not enable the design of fully deterministic algorithms to operate the management plane anymore. The approach here should be distributed and probabilistic management solutions. Their acceptance and deployment in very large networks remains, however, subject to several conditions. For every addressed problem and probabilistic approach provided, research remains to be done to:

- Clearly demonstrate the applicability of the probabilistic approach.
- Precisely evaluate the degree of uncertainty (the risk) a deployed probabilistic management approach can handle while operating safely.
- Demonstrate the gain for system/network administrators versus the overhead and risk.

Behavior of Managed Systems

Management is fundamentally about deciding and delivering behavior. We want to model and manage the behaviors of hardware, software, and even users within a system. Without the ability to make predictions about behavior, we cannot make service guarantees.

A lot of effort has been spent over the years on the construction of data models that represent the configurable attributes of devices. Data and information models, such as MIB and CIM, are large and elaborate models for registering the configurable parameters of existing hardware and software. However, if this effort has the aim of modeling behavior, the results have been far from successful. The most widespread of these, MIB modules, are used mainly for monitoring the state. If a device fails to respond to a query about its state, then no information about its current state or behavior can be decided.

The effort to model data is based on an unwritten assumption that configurations correspond to behaviors, that knowing the attributes that are programmed into a device is sufficient to learn what it will do (at some appropriate level of approximation). Unfortunately, this is incorrect except for the simplest automata.

Behavior implies the ability to predict changes in a system, either changes made autonomously or in response to input (events or programming). Behavior can be understood empirically or theoretically. The empirical study of behavior has been addressed mainly under the topic of anomaly detection.

- Detection of normal and then abnormal behavior (security, intrusion, resource failure, detection)
- Stabilizing behavior of nodes/systems

In the worst case, one has little information about what governs an entity, and one is reduced to observation, somewhat like zoologists observing animals in the jungle. In the best case, one has some knowledge of the inner workings and programming of entities and their environment, and one can make probabilistic predictions.

The current standard for system modeling in computer science is the unified modeling language (UML). UML models data objects, their relationships, their interactions, and what we would like to see of their internal behavior. Behavior, however, is that which is observed, not that which is planned, and results from the environment in which software is run. Recently *promise theory* has been proposed as a logical and probabilistic model for dealing with approximations to behavioral constraints. Promise theory takes the view that systems should be reduced to fundamental agents of change (not subjectively chosen *objects*) and that each component in a system specifies its individual behavioral constraints rather than its algorithmic details. One then uses the resulting network of behavioral promises to model the probable behavior of the collective.

Today we would like to express behavioral promises as Service Level Agreements (SLA) to root computing within a commercial enterprise. An important aspect of these agreements is what should happen if a party fails to deliver on its promised behavior. The agreement is used as a form of meta-level programming that deals with unreliability.

Intrinsically, behavior is something that is observed. A closer study of observed behavior in systems requires us to ask basic questions such as: what are the values we ought to measure to best understand particular behaviors? Which parameters dominate the behavior and when? For example, a routing table might determine the behavior of a router at a reasonable approximation when traffic is low, but when a network is saturated, it is the capacity resources that determine whether packets will be forwarded or mostly dropped.

Future research is required to investigate the relationship between behavior, economics, and uncertainty. Promise theory seems to be an interesting tool for this, because it easily incorporates all three and makes clear the connection to the theory of economic games. The actual relationship between observed system behavior and its configuration or programming is also a subject for further study.

Conclusion

This article discusses seven important challenges for research in the area of network management. The article reflects the outcome of a workshop

Management is fundamentally about deciding and delivering behavior. We want to model and manage the behaviors of hardware, software, and even users within a system. Without the ability to make predictions about behavior, we cannot make service guarantees.

Future research is required to investigate the relationship between behavior, economics, and uncertainty. Promise theory seems to be an interesting tool for this, because it easily incorporates all three and makes clear the connection to the theory of economic games.

that was organized jointly by the IRTF-NMRG and the EMANICS NoE that took place in October 2006. It should be noted, however, that the article does not claim to provide a complete overview of all possible research challenges in this area.

ACKNOWLEDGMENTS

The work reported in this article was supported by the EC IST-EMANICS Network of Excellence (#26854). We would like to thank all workshop attendees for their contribution and SURFnet for hosting this workshop.

REFERENCES

[1] Report of the NSF Workshop on Fundamental Research in Networking, Apr. 2003, Airlie House, VA; http://www.cs.virginia.edu/~jorg/workshop1/
[2] J. Schönwälder, "Overview of the 2002 IAB Network Management Workshop," RFC 3535, May 2003.
[3] A. Gonzalez Prieto and R. Stadler, "A-GAP: An Adaptive Protocol for Continuous Network Monitoring with Accuracy Objectives," *IEEE TNSM*, 2007, vol. 4, no. 1, June 2007.
[4] A. Pras and J. Schönwälder, "On the Difference between Information Models and Data Models," RFC 3444, Jan. 2003.
[5] A. K. Y. Wong *et al.*, "Ontology Mapping for the Interoperability Problem in Network Management," *IEEE JSAC*, vol. 23, issue 10, Oct. 2005.
[6] T. R. Gruber, "A Translation Approach to Portable Ontology Specifications," *Knowledge Acquisition*, vol. 5, issue 2, June 1993.
[7] K. Misra, *OSS for Telecom Networks, an Introduction to Network Management*, Springer, Aug. 2004.
[8] M. Waldburger and B. Stiller, "Regulatory Issues for Mobile Grid Computing in the European Union," *17th Euro. Regional ITS Conf. (Int'l. Telecommun. Soc.)*, Amsterdam, The Netherlands, Aug. 2006.
[9] R. Badonnel, R. State, and O. Festor, "Probabilistic Management of Ad Hoc Networks," *IEEE/IFIP NOMS 2006*.

BIOGRAPHIES

AIKO PRAS (a.pras@utwente.nl) received his Ph.D. from the University of Twente (UT) in 1995. The title of his thesis is *Network Management Architectures*. He is an associate professor at the DACS group of UT. His research interests include network and service management, with specialization in Internet management, traffic measurements, accounting, and security. He is an executive committee member of the EU FP6 EMANICS Network of Excellence (NoE).

JÜRGEN SCHÖNWÄLDER received his doctoral degree in 1996 from the Technical University Braunschweig, Germany. He is an associate professor of computer science at Jacobs University Bremen, Germany. His research interests are cooperative distributed systems, network management, wireless sensor networks, and network security. He is an active member of the IETF and chair of the Network Management Research Group (NMRG) of the IRTF.

MARK BURGESS received a Ph.D. in theoretical physics in Newcastle, for which he received the Runcorn Prize. He is a professor of network and system administration at Oslo University College. His current research interests include the behavior of computers as dynamic systems and applying ideas from physics to describe computer behavior. He is the author of the popular configuration management software package, cfengine.

OLIVIER FESTOR has a Ph.D. degree (1994) and a Habilitation degree (2001) from Henri-Poincaré University, Nancy. He is scientific leader of the MADYNES Research Team at LORIA-INRIA Lorraine. His research interests are in the design of algorithms and models for automated and scalable management for highly dynamic environments. He is project leader of the EU FP6 EMANICS NoE.

GREGORIO MARTINEZ PÉREZ (gregorio@dif.um.es) received his M.S. and Ph.D. degrees in computer science from the University of Murcia. He is an associate professor in the Department of Information and Communications Engineering of the University of Murcia.. His research interests include security and management of networks and services.

ROLF STADLER received a Ph.D. in computer science from the University of Zurich in 1990. He is a professor at the Royal Institute of Technology (KTH) in Stockholm, Sweden. His research interests include network and service management, specifically aspects of scalability and autonomic behavior. In 1992 he joined the Center for Telecommunications Research at Columbia University and later he became visiting professor at ETH Zurich.

BURKHARD STILLER received his doctoral degree from the University of Karlsruhe in 1994. Since September 2004, he is a full professor and the communications chair at UniZH, IFI. He was appointed assistant professor for communication systems at ETH in 1999 and held additionally a Chair at the University of Federal Armed Forces Munich (UniBwM), where he headed the Information Systems Laboratory (IIS).