Towards Management of the Future Internet

[1]Sung-Su Kim, [1]Young J. Won, [2]Mi-Jung Choi, [1]James Won-Ki Hong, and [3]John Strassner
[1]Dept. of Computer Science and Engineering, POSTECH, {kiss, yjwon, jwkhong}@postech.ac.kr
[2]Dept. of Computer Science, Kangwon National University, mjchoi@kangwon.ac.kr
[3]Waterford Institute of Technology, jstrassner@tssg.org

Abstract**—** While the current Internet is successful in many aspects, management of the current Internet has a set of associated systemic and business problems. This paper discusses those problems and presents a summary of management research efforts on the Future Internet. We present the design requirements for Future Internet management such as architecture, protocol, knowledge representation, and market aspects.

Index Terms**—Future Internet, Network Management, Manageability**

I. Introduction

The current Internet is one of the most successful technologies in history. It has become a social phenomenon, and has had a major impact on how we conduct business and changed the way that people communicate with each other.

However, the current Internet also has many problems associated with it, including security, mobility, QoS, scalability, and manageability. These problems have been dealt with in an ad-hoc manner quite successfully so far, but researchers have found that some of these problems come from fundamental problems with the original Internet design. Many researchers have started working on redesigning the Internet (so called Future Internet, Future Network or New Generation Network) in an attempt to solve these fundamental problems [4].

There are three fundamentally different approaches for the design of the Future Internet. The first is an incremental or evolutionary approach, and is evidenced by the plethora of solutions being applied to the current Internet that violate its architectural principles, such as Network Address Translators, firewalls, and Virtual Private Networks. [16] has documented many reasons why this approach is no longer sustainable. The second is what is called a 'clean slate' approach [25], which eliminates existing commitments, restraints, and assumptions, and starts with a new set of ideas. While the first approach evolves a system from one state to another incrementally, the latter enables a radical redesign of the current Internet architecture.

However, a third alternative exists, which is a compromise between the above two approaches. This approach enables new ideas to evolve while simultaneously emphasizing backwards compatibility with the existing Internet. This is very important to certain stakeholders, such as Internet service providers (ISPs), who have invested billions into their equipment and want to leverage those investments. It is motivated by the fact that current and future networks and networked applications have vastly different requirements; this means that a *single* architecture cannot simultaneously meet these different needs.

Each research community has its own focus on Future Internet design. Unfortunately, there is a tendency to ignore the management aspects of Future Internet. Some researchers consider the Simple Network Management Protocol (SNMP) [26] or current command line interfaces [31] appropriate for Future Internet management, but such network management protocols are not even suitable for the current Internet. For example, SNMP operates on top of the data plane and hence, management protocols rely on the correct operation of the very thing they are supposed to manage. Recently, [23] pointed out that the IETF had accepted the fact that SNMP and related approaches had failed. In our previous work, we summarized the management requirements of the Future Internet [15].

In this paper, we focus on the management of the Future Internet. However, this is not solely a technical endeavor; network management architecture and economic aspects should also be considered to develop new management technology. The rest of the paper is structured as follows. Section 2 investigates research efforts into managing the Future Internet. Section 3 describes various aspects of Future Internet management design. In Section 4, we draw our conclusions on the management of the Future Internet.

II. Related Work

The current Internet has shown that many problems originated from not considering management functionality during its design phase. Management features have been incrementally added to Internet and this has led to a number of problems, such as large complexity, scalability, and extensive manual involvement. Research on the management of the Future Internet, as well as evolutionary Next Generation Networking (NGN), is being performed worldwide; representative work in the EU, USA., and Korea will be summarized in this section.

A. Research in the EU

The AUTOI project is targeting evolutionary NGNs [10]. AUTOI rejects the current Internet's notion of a network that is service-agnostic, and instead envisions a set of networks that are directly aware of the services that they carry and hence, the requirements of each service. Each network will manage resources by applying autonomic principles [6] to develop a self-managing virtual resource overlay that can span heterogeneous networks and supports service mobility, security, QoS and reliability. In this overlay network, multiple virtual networks co-exist on top of a shared substrate with uniform control. The overlay will be self-managed based on the system's business goals, which drive the service specifications and implementations. AUTOI uses principles from the FOCALE [6, 11] architecture, which is a context-aware, policy-based autonomic architecture; this enables the services and resources that it offers to be dynamically adjusted according to changes in context. The AUTOI architecture is based on the concept of planes [10].

The goal of 4WARD project is to design and test a family of future global telecommunications networks (the 4WARD framework) that has the potential to supersede both current telecommunications networks as well as the current Internet [12]. Requirements from the main non-technical perspectives (usage and services, especially in an "Internet of [interconnected] Things"), socio-economics, regulation, governance, and policy) are used to derive the technical requirements of the 4WARD framework as well as its vertical technical themes (mobility, security, scalability, availability, QoS, inter-provider issues, and physical layer awareness). Special consideration is given to usability, sustainability, and environmental awareness. Note that 4WARD is not specifically chartered to address the design of "the" Future Internet; rather, it is viewed that such a design will emerge from considering other Future Networks that it must interact with. Traditional networks are managed by an external process; in contrast, 4WARD proposes a new paradigm for network management, where management functions come as embedded capabilities of the devices. With this approach, network elements have embedded management capabilities, consisting of several autonomous components which interact with each other in the same device and with other components in neighboring devices, enabling these network elements to form a thin "management plane" embedded in the network itself. This is different than the management approach of AUTOI, whose management and orchestrations planes are external to the devices.

ANA [13], like AUTOI, is focused on evolutionary NGNs. ANA first focuses on the so-called "waist" of the Internet: its IP layer. Variability in the current Internet occurs both above and below the IP layer, giving rise to the "hourglass" model. This is a fundamental reason why updating, let alone changing, the current Internet is so difficult. ANA thus focuses on disruptive technologies that "grow" a new architecture in two sets of design-prototype-test phases. ANA is not a "one size fits all" architecture, but rather a framework to host, interconnect, and federate multiple heterogeneous networks (which is why it is not a true clean-slate approach). This is done by using the ANA framework to specify how networks interact using a set of core abstractions.

B. *Research in the USA*

The FIND initiative [17] is supporting many research projects in the design of the Future Internet. FIND has two projects specifically related to the management of the Future Internet. The first project, Towards Complexity-Oblivious Network Management [14], proposes a new network management architecture based on an approach called Complexity-Oblivious Network Management (CONMan) [27]. In this architecture, the management plane is independent of the data plane. Management operations, which are carried out by software 'network managers' located on distributed network devices, use the generic management interface to configure and connect software modules on managed devices. This restricts the operational complexity of protocols to their implementation and allows the management plane to achieve high-level policies in a structured fashion. Policies themselves, defined and written by network administrators, are specified in a high-level, declarative, domain-specific language. This language specifies the general properties that are to be achieved and maintained; the translation from goal specification to actual implementation is performed by the software network managers. CONMan technologies will support other research efforts, such as Emulab and GENI.

The second project, "Design for Manageability in the Next Generation Internet (NGI)" [18], is focused on reliability and performance in the Future Internet. It assumes that most management functions will be automated and will rely on embedded capabilities in the network. Research is focused on exploring design alternatives for network management building blocks that will provide essential low-level functionality and can then be composed in different configurations to support management functionality. Results will include design documentation and in some cases, initial prototypes for various building blocks. The project also includes education and outreach activities that will develop course materials on network measurement, configuration and management.

The goal of the Operations, Management, Integration, and Security (OMIS) working group [20] is to specify requirements for operating GENI and managing its services both from a researcher's viewpoint and from that of its operators. They have already described the high-level functions required for coordinating operations in the near-term research and prototype environment and they will propose a long-term operations framework and define concepts of operations, including operations management.

C. *Research in Korea*

The Collect, Analyze, and Share for Future Internet (CASFI) [19] started in Korea in 2007 as an independent research group and was successful in soliciting a project funded by the Korean government for four years starting from March 2008. The authors of this paper from Korea are also members of the CASFI project. The current Internet carries many types of traffic, such as VoIP and IPTV, not originally envisioned at its inception, and offers many challenges in network performance monitoring and management. CASFI intends to develop sound and practical techniques and methods for performance measurement and analysis of the current Internet, and gain insight for better manageability in the Future Internet [24].

III. Design Requirements for Future Internet Management

A. *Information exchange between management layers*

The concept of layers is almost always used to build networks. Layers can be used as a powerful abstraction mechanism in software engineering, making complex problems much simpler to analyze, design and implement. In addition, layers have been used in the design of the current Internet. So, it sounds reasonable to use layers in Future Internet management architecture. However, management architectures need a richer, more flexible approach, since their goal is to integrate and harmonize data. This is especially true of autonomic management systems, which rely on semantic reasoning to make decisions. In other words, semantic reasoning infers logical consequences from a set of facts or axioms. As such, this is a search for meaning, and the fact that a fact or inference relates to a layer is irrelevant [2]. Thus, if the layer abstraction does not work, a new abstraction should be used. An alternative architecture that does not use layers is FOCALE, which is presented in Fig 1.

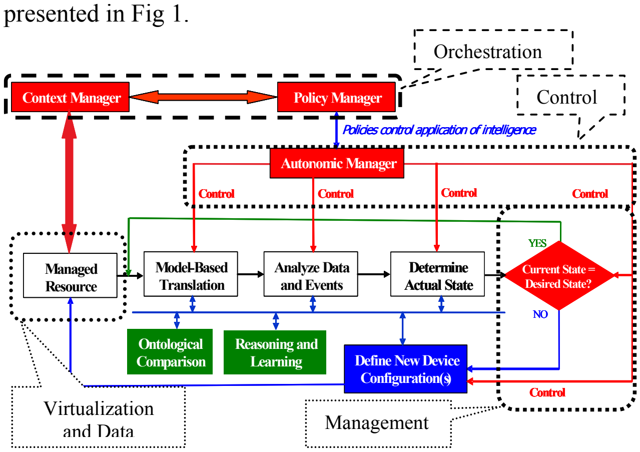

Fig. 1 FOCALE-based Version of AUTOI Architecture

The basic FOCALE architecture has not changed [11], because providing virtualization, management, and orchestration is already done without layers. FOCALE is a model-driven approach [28], and is based on the DEN-ng information model [29]. DEN-ng models management and orchestration, and FOCALE uses these models to orchestrate behavior using the above control loops. In addition, a virtual resource and a virtual service are already defined as managed entities, and virtual (as well as concrete) services are already defined as being hosted by virtual (and/or physical) resources. Hence, the managed resource could be a physical, virtual, or hybrid (e.g., a physical resource with virtual device or service support) managed entity.

Therefore, as we reviewed above cases, we can also manage the Future Internet without layers by using new architectures such as AUTOI and FOCALE.

B. New Protocol and Addressing Approaches

Current Internet addressing has several problems, such as routing scalability, mobility and multi-homing support. Better addressing approaches are required for the Future Internet.

The Public Switched Telephone Network (PSTN) network addressing separates users from the actual switching address. Therefore, the addressing scheme insulates the switching infrastructure from problems generated by users such as congestion and hacking. In contrast, the current Internet uses global IP addresses, which takes both identifier and locator roles. To reduce problems caused by this integrated addressing scheme, we need at least two or three separate address spaces.

An identifier has a role of making every object which has an identifier independent of topology and physical location. An identifier should provide the high-level functionality which will make the Future Internet extensible and useful for such things as access control, QoS provisioning, mobility and multi-homing support, dissemination, and content/service discovery. A locator should focus on fast and reliable transmission [1].

C. Separation of Data, Control, Management, and Inferencing Planes

In current data networks, the functionality that controls the network is divided into three main planes: the data plane that processes packets on a communications processor; the control plane that routes traffic on an application processor; and the management plane that monitors the network and configures the data-plane mechanisms and control-plane protocols. The original IP control plane was designed to be simple; the data plane is maintained by a single distributed algorithm. In contrast, the current data, control, and management plane is more complex than this original design concept, due to the proliferation of services that did not even exist when the Internet was first conceived. These new services have created many dependencies between the data and control planes, which are usually not automatically maintained [7].

The problem is that packet delivery cannot begin until the routing protocols create the necessary forwarding tables, and the management plane cannot reach the control plane until the routing protocols are configured. In most cases, to resolve this contradiction, configuration information must be installed on IP routers before the service can be deployed. This solution might require hundreds of thousands of lines of low-level configuration commands distributed across all relevant network devices. The result is a complex and failure-prone network.

The current state-of-the-art is focused on two things: optimizing the data plane (the data traffic represented as a path through the network) and applying mechanisms of the control plane (a set of discrete nodes that affect overall metrics (e.g., QoS) of the data traffic) as needed (but usually, on a point-by-point solution basis, especially since the current Internet is made up of heterogeneous devices having different functionality and no common programming language). The problem with this is that manageability is an afterthought, making it impossible to configure network services and resources to support the realization of business objectives, particularly across heterogeneous networks with different ownership and possibly conflicting objectives. Furthermore, different control plane mechanisms are used that cannot be easily coordinated.

[21] shows a conceptual overview of the inference plane to solve these and other problems. In reality, it consists of two separate planes, a management plane and the actual inference plane. The concept of “Control Plane” and “Management Plane”

has been fundamental to the telecommunication model. The control plane focuses on real-time signaling among network elements for service provisioning and traffic engineering, and implements distributed algorithms (e.g., routing) to enable data plane functionality. The management plane is a near-real-time approach for managing the network through appropriate external interfaces that the network devices provide. It is responsible for both configuring data plane mechanisms (e.g., addresses) as well as control plane protocols (e.g., routing policies). The inference plane *separates* this hidden dependency between the control and management planes in order to create independent, scalable governance functionality.

The inference plane is a coordinated set of intelligent decision-making components that represent the capabilities of the computing elements being controlled, the constraints placed upon (e.g., by regulatory and business policies [8]), and the context in which they are being used. Most importantly, it separates data, control, and management, enabling a better understanding of the overall environment to be constructed. This enables the services and resources supplied by heterogeneous networks to be assigned and driven by business goals and rules.

D. Monitoring and Measurement

Measuring and studying structure, performance and unacceptable behaviors of the Internet have received increased attention, as the Internet has evolved rapidly in the last decade. Measuring and understanding the Internet's properties is necessary to troubleshoot and debug failures faster and more efficiently, design future protocols that are better adapted to its special characteristics, monitor performance and react appropriately to change in real-time [9].

Measuring traffic in the current Internet is a daunting task, especially with an ever-increasing number of connecting nodes. Growth in the number of applications has also resulted in very complex traffic and usage dynamics. In the Future Internet, monitoring and measurement will be much more important, because management requirements will increase due to a larger number of different services being deployed. Management of the current Internet is concentrated on monitoring and measuring performance of network devices, not services or users. But, in order to provide various network services in the Future Internet, management of higher level objects (e.g., services and Service Level Agreements (SLAs)) instead of current lower level objects (e.g., device interfaces) is required. Managing services is becoming more and more complex and time consuming for service providers, since services are increasing both in number, complexity, and the number of users per service. This has created a desire to enable users to partly monitor and manage services of the Future Internet. This is called Customer-centric Service Management (CSM). The SLA plays a key role in CSM, as it can determine in detail the possible user service management actions.

E. Knowledge Representation and Reasoning

Currently, there are several integrated network management models using different technologies for resource management, such as SNMP (Simple Network Management Protocol), CMIP (Common Management Information Protocol), MIP (Management Information Format) and WBEM (Web Based Enterprise Management) [3].

Every integrated management model needs a specific management information definition language to describe the resources to be managed (its management domain), in order to communicate between managers and agents. Therefore, each model has defined its own management information definition language which has different capabilities and expressiveness:

- SMI (Structure of Management Information), with its different versions, for SNMP.
- GDMO (Guidelines for the Definition of Managed objects) for CMIP
- MIP (Management Information Format) for DMI
- MOF/CIM (Managed Object Format/Common Information Model) for WBEM
- IDL (Interface Definition Language) for CORBA.

Problems arise when different management models are used on the same networked system. Interoperability is necessary for the management models to provide a unified view of the whole managed system. However, this cannot be easily provided due to differences in syntax and semantics. The state of the art is to translate syntaxes; this means that when the same resource is described with two different management information definition languages, it is possible to apply a direct translation between the defined structures of the descriptions, but not between their meanings. Thus, in order to use different models, semantic interoperability is necessary. Unfortunately, all of the listed approaches are incapable of defining semantics, since they are not formal languages and thus do not support reasoning (using, for example, first order logic).

Interoperability of network management is not the only problem; there is the problem of knowledge representation and reasoning. SMI was built to define low-level concepts of network devices, such as a device interface. This is true of the other approaches as well. None provide the ability to understand management concepts, such as how to represent whether an SLA was violated or not. It is too difficult to modify existing approaches; therefore, we believe that a new approach must be created that can *leverage* existing low-level concepts and employ them in management scenarios.

Another problem is provability. Assume that a problem is found with a traffic scheduling algorithm on a certain router. It is not good enough to make a guess that will (for example) change the algorithm to a different type. We should be able to prove mathematically what, exactly, will happen when the algorithms and components are changed. Without management data based on formal logic, it is impossible to prove changes mathematically. There is an aspect of software engineering, called theorem provers [30], that show that a conjecture is a logical consequence of a set of statements. Our current research combines theorem provers with models (to represent network management *facts*) and ontologies (to formally define their meaning and semantics) in order to build a formal network management language [5].

In order to exchange and reuse common management data among different applications, a common information model must be used. Different applications will always use private object models and different repositories that require different data models. However, in order to share and reuse data, a common set of mappings that use the same information model to derive each data model must be used. This enables the development of mediation software that can translate different representations of the same data into different repositories, as well as translate queries that take different forms because different protocols are used by different repositories. We use the DEN-ng model for this purpose, as it is the only model that is inherently extensible (due to its being based on software patterns), contains important sub-models (e.g., context, metadata, and policy), and was built to support ontologies and formal management approaches (e.g., based on finite state automata).

F. Decisions using Uncertain or Incomplete Data

Often, network management data is incomplete, due to the data being contained in multiple non-discoverable locations or simply through transmission error. Another source of problems is that the system may not be certain as to the reliability and/or correctness of data. One way to solve these problems is to convert a single *correct* definition of system behavior into a set of *acceptable* system behaviors. Then, the manager selects *one* behavior to use. The use of a formal language to describe these behaviors is critical, since this enables their systematic fuzzification and facilitates graded conclusions based on probability and other mathematical concepts.

G. Managing Nodes based on System and Application Needs

The current Internet routing system is based on global IP addresses, and traffic is transmitted by a simple 'best effort' principle. The current Internet does not have the concept of differentiated traffic treatment. Large complex services or systems of the Future Internet, such as enterprise websites or distributed information systems, will be composed of both interacting and independent components. To reflect system and application needs, component based approaches should be considered. In order to meet SLAs, many systems have used over-provisioning [22]. A typical example is searching, which uses extensive over-capacity to guarantee time limits on search requests.

Since there will be more and more applications in the Future Internet, and network resources are limited, over-provisioning is not appropriate. Therefore, we need more efficient management of network resources. It is hard to precisely characterize or limit the dynamic behaviors of the composed enterprise systems and applications. This characteristic makes it difficult for system administrators to efficiently achieve conformance to SLAs. System's constituent components should be individually monitored and managed to the extent needed to attain desired runtime component behaviors. In order to monitor each component by its own objective, mechanisms to match high level goals to component level objectives should be developed [6, 8].

H. Market Aspects

The management of the Future Internet must also consider economic aspects.

The first aspect is reputation of management technology. Assume that we choose SNMP as a Future Internet management technology. This may complicate the tasks of network administrators and Internet network operators, due to the inherent problems of SNMP in security, complexity, speed, and limited functionality.

The second aspect is market control and protection. Every company wants to secure a lucrative niche market; management technologies can be included as well. Some companies contribute to standards because they want to keep control over key technologies and their associated markets.

The third aspect is that currently, a sustainable market for open management does not exist, due to the problems already cited in this paper and in [23].

Considering these three aspects, Future Internet management technology should prove that it can be profitable to vendors, administrators, and other stakeholders.

I. Automation of Management Functionality

At the same time that networking technologies and capabilities are expanding, there is an ever increasing need for streamlining and economizing network operations. Network management teams in the Future Internet must be able to operate and maintain their infrastructures with the same or fewer human resources than today's networks [18]. The motivation behind autonomic management is to identify those functions that can be done without human intervention to reduce the dependence on skilled resources for managing devices, networks, and networked applications. Benefits include reduction of manual errors and faster configuration, which means faster service billing. If the autonomic network can perform manual, time-consuming tasks (such as configuration management) on behalf of the network administrator, then that will free up the system and the administrator to work together to perform higher-level cognitive functions, such as planning and optimization of the network. This holism is fundamental to solving system-of-systems problems [24].

IV. Conclusion and Future Work

We have summarized the state of worldwide research in the management of the Future Internet and evolutionary NGNs. Representative research was presented from the EU, USA, and Korea. We discussed design aspects of the management of the Future Internet, including management layers, new protocols and addressing approaches, monitoring and measurement, the separation of data, control, and management planes, knowledge representation and reasoning, and market aspects. We found that there are more complex requirements for the Future Internet management compared to the current Internet.

Our future work will focus on developing management information modeling, which will formalize the nature of data and semantics.

Acknowledgments

This work was partly supported by the IT R&D program of MKE/IITA (A1100-0801-2758, "CASFI : High-Precision Measurement and Analysis Research"] and by WCU (World Class University) program through the Korea Science and Engineering Foundation funded by the Ministry of Education, Science and Technology (Project No. R31-2008-000-10100-0).

References

[1] J. Choi, C. Park, H. Jung, T. Kwon, and Y. Choi, "Addressing in Future Internet: Problems, Issues, and Approaches," 3rd International Conference on Future Internet Technologies, June 2008.

[2] J. Strassner, S. van der Meer, and James W. Hong, "Autonomic Management of Communications Networks without Using the 'L' Word," submitted to JSAC, 2008.

[3] D. Clark, C. Partridge, J. C. Ramming, and J. T. Wroclawski, "A knowledge plane for the internet," In Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols For Computer Communications (Karlsruhe, Germany, August 25 - 29, 2003). SIGCOMM '03. ACM, New York, NY, 3-10.

[4] T. Roscoe, "The end of internet architecture", Proceedings of the 5th Workshop on Hot Topics in Networking (HotNets-V), Irvine, CA, USA, November 2006.

[5] J. Vergara, V. A. Villagra, J. I. Asensio, and J. Berrocal, "Ontologies: Giving Semantics to Network Management Models", IEEE Network, special issue on Network Management, Vol. 17, No. 3, May/June 2003.

[6] J. Strassner, "Autonomic Networking: Theory and Practice", Tutorial, 2008 IEEE Network Operations and Management Symposium, April 2008.

[7] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "A Clean Slate 4D Approach to Network Control and Management", ACM SIGCOMM Computer Communication Review. 35(5). October, 2005.

[8] J. Strassner, "DEN-ng: achieving business-driven network management", Network Operations and Management Symposium, 2002. NOMS 2002. IEEE/IFIP , pp. 753-766, 2002.

[9] T. Spyropoulos, S. Fdida, and S. Kirkpatrick, "Future Internet: Fundamentals and Measurement", ACM SIGCOMM Communication Review, volume 37, April 2007.

[10] Autonomic Internet (AutoI), http://ist-autoi.eu/autoi/.

[11] J. Strassner, N. Agoulmine, and E. Lehtihet, "FOCALE – A Novel Autonomic Networking Architecture", ITSSA Journal, Vol. 3, No. 1, May 2007, pp 64-79.

[12] 4Ward, http://www.4ward-project.eu/.

[13] Autonomic Network Architecture (ANA), http://www.ana-project.org/.

[14] F. Paul and L. Jay, "Towards Complexity-Oblivious Network Management", FIND project, http://www.nets-find.net/Funded/TowardsComplexity.php.

[15] Sung-Su Kim, Mi-Jung Choi, and James W. Hong, "Management requirements and operations of Future Internet", 11th Asia-Pacific Network Operations and Management Symposium (APNOMS 2008), LNCS 5297, Beijing, China, October 2008, pp. 156-166.

[16] D. Clark, K. Sollins, J. Wroclawski, D. Katabi, J. Kulik, X. Yang, R. Braden, T. Faber, A. Falk, V. Pingali, M. Handley, and N. Chiappa, "NewArch: Future Generation Internet Architecture", NewArch Final Technical Report, http://www.isi.edu/newarch/.

[17] Future Internet Network Design (FIND), http://find.isi.edu.

[18] B. Paul, B Suman, and E. Cristian, "Design for Manageability in the Next Generation Internet," FIND project, http://www.nets-find.net/Funded/Manageability.php.

[19] Collect, Analyze, and Share for Future Internet (CASFI), http://casfi.kaist.ac.kr/.

[20] Global Environment for Network Innovations (GENI), http://www.geni.net/wg/omis-wg.html.

[21] J. Strassner, M. Foghlú, W. Donnelly, and N. Agoulmine, "Beyond the Knowledge Plane: An Inference Plane to Support the Next Generation Internet", IEEE Global Information Infrastructure Symposium (GIIS 2007), 2-6 July, 2007, pages 112-119.

[22] V. Kumar, K. Schwan, S. Iyer, Y. Chen, and A. Sahai, "A state-space approach to SLA based management", IEEE Network Operations and Management Symposium (NOMS 2008), April 2008, pp.192-199.

[23] J. Schonwalder, A. Pras, and J.-P. Martin-Flatin, "On the future of Internet management technologies", IEEE Communications Magazine, vol.41, no.10, pp. 90-97, October 2003.

[24] Mi-Jung Choi, Sung-Su Kim, James W. Hong, and J. Strassner, "Towards the Future Internet and Its Manageability", Submitted to IEEE Communications Magazine, October 2008.

[25] N. McKeown and B. Girod, "Clean slate design for the internet." Whitepaper; Information available at: http://cleanslate.stanford.edu/, April 2006.

[26] D. Harrington, R. Preshun, and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol Management Frameworks", RFC3411, STD0062, December 2002.

[27] H. Ballani and P. Francis, "CONMan: A Step Towards Network Manageability", ACM SIGCOMM Computer Communication Review, Vol. 37, Issue 4, October 2007, pp. 205-216.

[28] www.omg.org/mda.

[29] J. Strassner, "DEN-ng Model Overview", Joint ACF, EMANICS, and AutoI Workshop on Autonomic Management in the Future Internet, May 14, 2008.

[30] M. Fitting, "First Order Logic and Automated Theorem Proving", Springer-Verlag, second edition, 1996, ISBN 0387945938.

[31] http://www.cisco.com/warp/cpropub/45/tutorial.htm, accessed February 6, 2009, from this URL.