

An Architecture for Network Management

Karen R. Sollins

MIT Computer Science and Artificial Intelligence Laboratory
32 Vassar Street
Cambridge, MA, USA
+1 617 253 6006
sollins@csail.mit.edu

ABSTRACT

As is becoming increasingly understood, in extending the Internet architecture into the future, network management is a key challenge. [4], [25] The current approach has been to provide a set of weakly integrated tools to network managers of each enterprise or other network. In this paper we argue the position that a single architecture or framework for network management would improve our overall ability to manage networks in an increasingly integrated, heterogeneous, and widely distributed network environment. We add to this the problem that increasingly the users and other components of our applications are distributed and/or mobile. In this paper we argue for such a common approach, and propose a set of key elements of such an approach, as a proof of concept argument that it is possible.

Categories and Subject Descriptors

C 2.1 [Computer Systems Organization]: Computer-Communications Networks, *Network Architecture and Design, Distributed Networks, Internet*. C 2.3 [Computer Systems Organization]: Computer-Communications Networks, *Network Operations, Network Management*

General Terms

Management, measurement, performance

Keywords

Network management, knowledge plane, information plane, regions.

1. INTRODUCTION

The focus of this position paper is that a common set of architectural design principles and framework for network management will allow for a significant improvement in being able to provide scalable, adaptable, distributed network management. The objective is to meet the network management requirements for our growing and evolving network environment.

We begin the paper with three brief examples, in order to demonstrate both the breadth of functionality of network management tasks and the systems challenges they face. The three examples do not cover the full breadth of network management problems, but are intended to provide the reader with a sense of the breadth of challenges that fall into scope for network management.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ReArch '09, December 1, 2009, Rome, Italy.

Copyright 2009 ACM 978-1-60558-749-3/09/12...\$10.00.

The examples are that of a mobile employee outside the home enterprise, the health of distributed applications, and zero-day anomaly detection.

In our mobile employee example, the intention is that as much as possible, the employee's experience should be the same inside and outside the home enterprise. Consider the situation in which the mobile employee is collaborating with employees inside the enterprise. There are two key dimensions to the problem. The first is that from a security, authentication and authorization perspective the mobile employee should be able to function effectively. The NAP/NAC architecture [3] developed by Cisco and Microsoft provides an approach to verification, authentication and authorization that involves several border functions, verification and, when that fails, conformance, in order that it succeed. Because this can be based on end-to-end algorithms, the communication can be isolated from any underlying potentially untrustworthy components. The second dimension to the problem has to do with resources and performance, and must address whether and how the communications resources (bandwidth, latency, reliability, etc.) can be provided or if not, the problems analyzed and mitigated. In this case, there is no isolation from the underlying resources, and the network management team on which the mobile employee is dependent does not even have control or possibly access to those other intermediate resources. The problems faced in this case are not only that of collaborative or cooperative network management, but also proprietary issues and scaling, under unpredictable conditions, because the mobile employee's location may be both new and changing.

In our second example, we consider a distributed application in which the distributed components may be running in different network management domains. We further envision that some or all of the components may be movable in order to address performance and other requirements. It would be valuable to be able to state and have applied performance criteria of the underlying network. In the case of some applications, they cannot run effectively if they do not have access to appropriate resources, or may simply run more effectively if they have more resources. This requires much of the same kind of analysis of conditions as that for diagnosis of failures or unacceptable behaviors, although in one way significantly different. Application health may require prediction from the supporting infrastructure, reflecting what will be possible in the future, as well as performance tuning when the predictions are inadequate. This may be important for an application that must continue to be active or available at some level of performance over some extended period of time. Network management activities may require not only multi-domain coordination, but also joint prediction.

In our third example, we consider the discovery of low volume zero day attacks or anomalies [23]. One aspect of this problem space is that there is no a priori signature or model of the offending traffic. It is unknown at least until it appears. This means that the first location of identification of the traffic must be at the destination points,

because only there can one distinguish wanted from unwanted traffic. A second is that if the traffic is low volume, the evidence will be sparse. The coincidence of these problems suggests that only through statistical aggregation across (weak) local detectors might one build more confidence in detection of such traffic. A similar approach is taking by Huang et al. [14] to diagnosing network disruptions. There are several cross-domain problems that arise in such situations. The first is that one needs to select the statistical aggregation algorithms appropriate to the underlying conditions.¹ In addition to the issues of coordinating across administrative boundaries, one must also organize the aggregation activities with the underlying models in mind.

We have touched here on several significantly different kinds of activities that fall under the umbrella of network management, e.g. disruption diagnosis and repair, prediction, and performance improvements. The reason for proposing a framework or architecture for network management is two-fold. The first is that the network management problems pervade and transcend the many more local network management domains in existence, and in some case, little or nothing exists at present. This includes integrating across the multiplicity of technologies that may support any communication path. The second is that there are opportunities for complementary and composite activities that otherwise are either unavailable, or lead to duplication of those activities.

This paper proceeds by identifying the key challenges for such a network management architecture, followed by the core elements of our proposed architecture, a brief summary of related work, and finally discussion of some of the particularly challenging open research problems in this area. This paper is intended to provide initial structure to a discussion in the research community about whether, how and why developments in network architecture can and should include network management and manageability.

2. THE KEY CHALLENGES

At its heart, network management is about the performance and behavior of the networks used for transporting communications. This may take the form of forensics after problems have arisen, health of the networks in order to avoid undesirable behaviors, improvements in order to increase the quality of the behaviors, and prediction of network behaviors into the future. Currently, network management might be considered to be a “cottage craft”, handled by a small cadre of extremely expert network managers, each working in a local domain, collaborating as needed on an individual and personalized and often ad hoc basis. The central challenge derives from two facts. First, there are large regions of “the network” which are unmanaged and these are likely to become increasing prevalent. Second, human-based and human-scaled peering for collaboration does not scale to the size, scope and complexity of the internetworking of the Internet. The challenge of this paper and of network management broadly is to provide a network management underpinning integrated into a network architecture (either current or future) that provides a reasonable level and breadth of network management and supports the expertise of the cadre of network managers with whom must remain many of the decisions about objectives, requirements, and constraints, both physical and policy

¹ For instance, one needs to understand the dependence or independence of the local weak detectors. In our prior work, we used SHT with an assumption of dependence in some places and HMMs where one can assume independence.

based. This is a call for support for network managers, not replacement of them.

To delve more deeply, we divide the architectural challenges to network management into two general areas, information and the reasoning and computation over it. Before discussing the challenges to each of these aspects of network management, we must first layout several key underlying assumptions, which in turn influence how we view the challenges in each of our two areas. The primary assumptions are: (1) the extent of network managers’ control is limited to their own local network; (2) those administrative domains also reflect proprietary and other policy boundaries; (3) need for network management crossing and extending beyond individual administrative domains; and (4) efficiency and performance, both because responses are often required quickly and because any activities should have minimal impact on the core transport service of the networks involved. The first three of these are probably true of any distributed capabilities. The fourth is reflective of the fact that network management is an ancillary capability, rather than primary to the purpose of the networks involved. With these in mind, we can now consider our two key aspects of network management, the information and the reasoning and computation over that information.

2.1 Information

Without information there cannot be any diagnosis of problems, optimization or improvement of resource usage, health analysis, or prediction more generally. Information may come from many sources, such as self-monitoring of resources, measurement tools, inference tools, and other network management tools themselves. Each type of information collection may run continuously, regularly, or only on demand. Broadly for network management, the collection and generation of information presents us with a number of challenges or requirements. We identify seven key ones here.

1. *Storage*: Storage is central to the collection of information, if the information is to be used after the fact. Often, even for immediate use, there is more information that can be collected in memory. Because storage often requires resources (energy, space) not available at the location of the information generation, at least some storage may be remote. A significant part of the storage challenge is to make enough of it available as needed. Key problems related to placement include little or no local storage, no single physical location of an artifact, and the need for non-local information, for example for forensics. The problems of storage are closely related to both performance and policy control below.
2. *Discovery*: One of the first steps in network management is to determine whether the information required is available. Since the analysis may involve multiple networks including transit nets, being able to learn what information is or is not available is key to any analysis. Note that one aspect of discovery may be to understand where inferences over existing information may be an adequate substitute for direct data collection or other measurement.
3. *Find information*: Simply knowing that the information exists or can be inferred is not enough. One must also be able to access it. This may mean moving a copy of the information to somewhere more immediate or moving the computation or query to the information itself. This challenge is also closely related to those of policy and composition of information.
4. *Share information*: A common, but inefficient, approach to network management is to collect the same information for each stove-piped network management tool. To the extent that

information collection uses network, computation or storage resources, minimizing the amount of duplication is advantageous. Sharing information also allows for subdividing the total work to be done, in order to make it most efficient and effective. This leads to a requirement for sharing information. Again, this challenge will need to be integrated with the issues of policy control.

5. *Reason over information*: If information is to be stored and shared, it becomes critical to understand it, independently of the measurement or analysis tools that may have generated it. This is true not only for individual information items, but also for the relationships among them. It is only with such a specification that rational reasoning over the information can occur, outside the information's original context. This is important for forensics, prediction, long-range analysis, and so forth.
6. *Extensible life and location*: Closely related to storage and the ability to reason over information is the requirement that at least some information have value over an extensive lifetime. During that time it may be aggregated, summarized, moved, reorganized and so forth. This challenge is also closely related to the storage, discovery, location, and policy challenges.
7. *Policy formation/composition*: The policy area is particularly complex. Policies are intended to enable controlled and limited access on some basis (economic, legal, organizational, or other privacy criteria). In the current mode of operation, access is provided often in real time and under duress, but with little formality to policies. Most often, network managers will officially take the position that no information is sharable or exposable. In practice², sharing of information does often occur, but often spontaneously. We suggest here that being able to formalize such control of access, especially in a context sensitive way, would be valuable, and lead to an improved understanding of what information is and is not available. A more complex challenge in this area derives from the fact that in many cases, the information required is not directly available, but can only be inferred from computation over sets of information. When those results are produced, one must derive or specify some composite policy. How to do that composition is not well understood, and may not be generalizable. Without an explicit specification of base level policies such composition is impossible. One level removed from the actual policies, it will also be critical for decision-makers to evaluate the tradeoffs in making information available under specified conditions. Such evaluation may be required not only based on monetary value, but perhaps risk, trust, and possibly social capital. An incentive evaluation component will play an important role.

Broadly speaking, the challenges or requirements for information fall into the problems of collecting, storing, using, and controlling information in a distributed, long-lived context, in which collection and usage of that information is only a secondary objective and hence must have little impact on the primary objective of providing transport. That said, information is at the heart of network management.

2.2 Reasoning and computation

In addition to the challenges and requirements of information, we are presented with a comparable set of challenges to the inference, analysis, and reasoning, i.e. the computational aspect, of network

² Based on discussions in the most recent NANOG and the GENI Measurement Workshop, Madison, WI, June, 2009.

management. Again, the challenges arise from the tussle among computation, performance, distribution, scaling, interests or objectives. At its heart we find the local network manager who needs to be addressing distributed problems, but whose first instinct, for a variety of scaling, proprietary and other reasons, is to retain as much control as possible, both of information and analysis. When we tease this apart, our current list of more specific challenges is:

1. *The nature of the information*: The information available for network management generally is incomplete, often incorrect, or intentionally masked or hidden.³ By definition the best one can expect is that it is statistically representative. The computations performed over the information must not only account for this statistical, incomplete, and large-scale nature, but must also be selected to match characteristics such as dependence or independence, long-term drift, availability or lack thereof of adequate training sets, and so forth. Generally, the appropriate algorithms are most likely to come from statistical machine learning, but must be selected appropriately to the intended results.
2. *Efficiency/performance*: As with the measurement, monitoring, and management of information, the tools that comprise network management can themselves also have an impact on the efficiency and performance of the network itself. Algorithms and their implementations must be selected to minimize their impact on performance while providing adequate functionality, accuracy, and detail, but need not provide more than that.
3. *Decomposition*: In many cases, it will be critically important to be able to decompose and distribute the computations required for a tool or capability. This may be driven by performance or policy constraints on moving information, competition for computing resources, or simply required computational capabilities, such as the limited availability of a particular machine architecture that especially suits a particular functional component. The drivers for decomposition include functionality, geography or topology, and policy.
4. *Composition*: As with the challenge of sharing and reusing information, it will be important to be able to compose tools or computations into more sophisticated ones. At a basic level, there are and will continue to be a set of tools that do inference over collected data. A simple example is the inference of lost packet rates. In fact, this may be handled by different tools in different places, but being able to compose a tool for aggregating traffic loss over the links of a path that traverses many networks will require the ability to compose not only the inferences of traffic loss based on local packet traces, but also tools for discovering the links that comprise a path.⁴ For broad network management capabilities to operate usefully and effectively it is important to support composition of functions or tools.
5. *Extensibility*: Beyond the problem of composition, we must recognize that network management is not static. New tools and capabilities are being designed and implemented all the time. For network management not to stagnate and become out of date, it is necessary that a composed tool be extensible to incorporate newer and more effective supporting tools as they become available. The challenges here include not only how to

³ Again, our conversations in the GENI Measurement Workshop and at NANOG, as well as personal experience, confirm this.

⁴ This is intentionally overly simplistic for purposes of the example.

discover them, but also how to invoke or use them, without requiring human intervention for each new extension.

6. *Organizing framework*: In order to relieve network managers of the often extremely detailed and complex tasks of organizing network management functions under differing and often changing conditions of physical organization, behavior criteria, and policy constraints, there is a need for a framework that can operate over a set of constraints and objectives to organize network management tools appropriately and then evaluate the effectiveness of that organization. Even the initial structuring, organizing and location of functionality is not an easy task. There is a significant open question of how to be able to evaluate the effectiveness of a framework set of decisions for organizing network management function. One must ask about a variety of questions such as the nature and source of the evaluation criteria such as what are they, how accurate or flexible can they be, etc. In addition one must ask about the scope and time over which such evaluations might take place.

3. ELEMENTS OF THE ARCHITECTURE

Our proposed architectural approach to this set of network management challenges is to expand on the idea of the Knowledge Plane (KP)[5]. The vision of the KP was to make the network(s) self-knowledgeable, self-analyzing, self-diagnosing, self-optimizing, and generally self-managing. We have come to refine that with the realization that it is unlikely that we either can or want to get the humans completely out of the loop, but there is a long way that we can and should go to providing automated management support and knowledge integrated into the network architecture. The architecture as we summarize it below falls into two primary components, the underlying information plane and the reasoning framework. We will address each separately.

3.1 The information plane

At the core of the information plane are information items. The sources of these items include measurement or other acquisition tools, factoids that may for example be written down or learned from such things as configuration files, etc., and the results of computational tools that execute in the KP. Each information item has an identifier. An information item can contain as part or all of its content other information item identifiers, facilitating complex items. An information item will be an instance of a class or type. Hence there is need of a type or ontology system. The class will provide both an abstract structure for the item and the semantics of that item. An information item may also have associated with it meta-data. Examples of meta-data include location criteria such as where and where not it may reside, ownership, access control policies, creation information, perhaps size or other aspects that may have an impact on performance, and so forth.⁵

In order to address issues of scaling, whether in physical or topological space, or more abstractly in terms of interest, we identify the need for a special kind of information item, the *region*. A region is a grouping of other information items, and as such provides a means of partitioning the universe either for straightforward scaling or for functional or policy boundaries. Because of their nature of grouping sets of identifiers and their meta-data, regions also provide

⁵ We recognize but do not discuss here the fact that ownership, policies, etc. require some concept of principals. These can be mapped onto an underlying information centric universe, but is left to a longer paper.

the basis for a publish-subscribe paradigm for producing and consuming information. Regions will also organize, so that among sets of them, they will share advertisements of the information items they know about, expanding but controlling the horizon of an information item.

The design of the information plane is based on several related efforts. First, in our earlier *Information Mesh* Project [30] we developed experience and a deeper understanding of the value of globally unique identifiers for information in order to support scaling and longevity. More recently, in Lee's dissertation [22] we worked with the ontology language OWL [33] as a starting point for some aspects of the KP (an alternative here might be CIM [8], but choosing a single starting place is key), and in Li's dissertation [23] and in Sollins' paper [29] we clarified the need, definition and usage opportunities of the *region* concept. We are currently beginning to design on top of an information centric substratesuch as that available from the PSIRP project [32]. PSIRP provides identification, publication, subscription and transport of information items, with no reference to ontology or higher level functionality. It also supports line speed policy based routing at least in a local context, based on a Bloom-filter approach[17]. Clearly, we will need to enrich such a model for our needs in the information plane, but it provides a strong starting point.

3.2 The Knowledge Plane: reasoning

The second aspect of the architecture is what takes it from being information to a knowledge plane. This is the capability of reasoning over the information to understand, hypothesize, infer, and act on knowledge rather than basic information. The knowledge aspect of the framework will include in its core:

- *Ontology*: An ontology not only allows for syntax and semantics of the information, but enables or constrains the scope of reasoning that can be performed on the entities defined by the ontology. An ontology in the KP must support extensibility, locally independent definition, some reasonable amount of convergence, and global discovery when needed. Our current ontology language of choice is OWL [33] although it is not the only possibility.
- *Function library and definitions*: There are two reasons that a library or catalog of network management tool definitions and implementations is important to the architecture. First, because the management target is any part of the broad network where management is desired, tools may be needed in a wide variety of locations. Perhaps even more importantly, in order for the KP to improve and evolve, it will be important to incorporate new tools with new capabilities into existing toolkits. This will require both a definition of each tool and implementations. If each inclusion of a new capability needs to be handled through manual intervention, improved and evolving behaviors are unlikely to succeed.
- *Probabilistic programming*: The significant majority of computation that will occur in the KP will be statistical or probabilistic. Lee [22] in his thesis took a preliminary step in specifying probabilistic knowledge. Beverly [1] in his thesis concentrated exclusively on statistical analysis of network information, because essentially all information that is collected from measuring and monitoring is sampled, incomplete, only partially accessible, intentionally incorrect, or some combination of these. The information as a whole is statistical.

- *Agent system:* As we have discussed above, network management must be partitionable across a variety of locations for reasons ranging from technical such as performance and efficiency based on physical distribution to policy such as information or tools being proprietary or under other policy constraints. Therefore many information gathering and information processing activities will by necessity be distributed, leading to partitioned functionality and a requirement for coordination and cooperation. For this, a basic agent system that supports such partitioning, communication, and possibly mobility will be required. In general it is a challenge because designing a secure and trustworthy agent system has remained a challenge.
- *Reasoning organization framework:* One of the core challenges in a design that requires as much distribution, coordination, extensibility, and policy control as the KP is that the questions of (1) how to decompose functionality in order to distribute it, (2) how to re-organize functionality under changing conditions, and (3) how to understand the effectiveness of an organization or re-organization of functionality, will require at least automated assistance for the human programmer or network manager. We suspect that the integration across many factors, locations, and policies is not something that human intelligence is best suited for. The humans will definitely be the sources of policy definitions and choices. They may oversee and supervise the organization and distribution of tool functionality, where more abstract and higher level reasoning is critical, but systematic and large volume-information-based decisions and reasoning are probably best handled through automation. Thus, the KP requires an organizing framework that can evaluate performance and distribution constraints, functional decomposition, and policy constraints in order to organize the composition of functional subcomponents, and then monitor and possibly revise such a structure under behavioral expectations.

3.3 Related work

Due to lack of space, we will only very briefly review a small sampling of related work and topics of related work.

We begin with the overall design and approach to network management. As mentioned earlier, the originally idea of an architectural perspective on network management derives from Clark et al.'s paper on the Knowledge Plane idea[5]. More recently Greenberg et al. described their 4D approach to route table computation, beginning with offloading it from routers and then proposing an AS wide computation. [12]. This approach is a first step in identifying and understanding the distinct challenges to information and computation. More recently, Li [23] in his work on zero-day worm analysis and Huang et al. [14] both proposed a computation aggregation approach to addressing distributed network management problems. Lee [22] proposed a richer compositional model for diagnosing problems. These only scratch the surface, but provide a sampling of related work.

Another thrust that is related to ours is the vision of "autonomic networking" or "autonomic computing". See [9], [15], [18], [19] for selected references. The source of this idea was IBM, in order to address management problems originally in the context of large scale distributed computing facilities, which were becoming unmanageable by humans. Two key components of their model are the autonomic control loop and the knowledge at the core of the

system, whether a distributed computing system or a network. Here we address the architectural drivers and criteria for such an approach and demand that it function in a much more diverse and tussle-filled environment.

With respect to a common substrate for information, one of the earliest proposals was Sophia [34] in PlanetLab. More recent further work in PlanetLab includes the PCL database [27], CoMon [26], and other PlanetLab wide facilities, e.g. [2], and [13]. These provide a hint to us, but PlanetLab does have not of the scale or distributed concern challenges that are faced more broadly in the Internet, nor decentralized management or control of management. The iPlane project [24] collects network wide information on path behaviors, in order to support some of the needs of overlay and peer-to-peer systems. More generally, in line with our information based model of the information plane, prior work derives from general pub/sub systems [10] and approaches to information or content centric networking such as [16], [32], [7], [21].

A third area of significant prior work is in the statistical approaches suited to the nature of the information and goals in network management. Huang et al. [14] and Li [23] demonstrate particular algorithmic choices. Goodman et al. [11] and Phillips [28] define statistical programming languages, and Lee [22] takes a first step at introducing the idea of statistical or probabilistic information.

We only mention very briefly here a fourth area of related work, agent systems. This is a huge area to which we cannot possibly do justice, but point out two projects that have influenced our work that of Tripathi in his Ajanta Project [31] and Li's dissertation [23]. The key for this work is not the agent system itself, but organizing it. Sollins' work on regions [29] and Li's dissertation provide the basis for much of that thinking. Related to this concept is the idea of *scopes* from the PSIRP project. [32]

4. CONCLUSION: KEY HARD RESEARCH PROBLEMS

In considering our architectural approach to network management, at this point, we suggest four key research problem areas here.

How to understand organizational constraints: Organizing and managing the potentially distributed information and knowledge planes of this effort must take into consideration such issues as physical location, topology, political location, functional (de)composition, and legal, economic and social policy constraints during frequent or constant changes to any or all of these. Composing these constraints is not well understood.

Impact on the network: Network management must always take into account that it is not the primary task of our networks, but only secondary to actual transport. Being able to evaluate and model such impact is important.

Managing information: Several key challenges in managing information have to do with appropriate storage and caching, abstracting and summarizing, and a model of lifetime of information.

Tussles: A system such as this is rife with tussles [6], differences of opinion or concerns. Finding a way to identify, express, and resolve those will enhance such a system significantly.

To conclude, we propose that an architecture for network management is valuable, feasible, and provides a number of open research problems.

5. ACKNOWLEDGMENT

This work has been supported by grants from The National Science Foundation (#0137403 and #0626904), Cisco Systems and Intel Corp. Thanks also to the three PhD students on this project, but are currently unavailable for joint authorship: Rob Beverly (BBN), George Lee (Google) and Ji Li (Credit Suisse).

6. REFERENCES

- [1] R. Beverly, "Statistical Learning in Network Architecture", PhD Dissertation, MIT Department of Electrical Engineering and Computer Science, May 2008.
- [2] P. Brett, M. Bowman, J. Sedayao, R. Adams, R. Knauerhase, A. Klingaman, "Security the PlanetLab Distributed Testbed", LISA '04, 2004.
- [3] Cisco Systems and Microsoft Corporation, Cisco Network Admission Control and Microsoft Network Access Protection Interoperability Architecture, white paper, Sept. 2005
- [4] D. Clark, "What is a Future Internet Architecture", Future Internet Information Meeting, Arlington, VA, July 29, 2009, slides and talk available at: <http://www.visualwebcaster.com/event.asp?id=60947>.
- [5] D. Clark, C. Ramming, C. Partridge, J. Wroclawski, "A Knowledge Plane for the Internet", SIGCOMM 2003.
- [6] D. Clark, J. Wroclawski, K. Sollins, B. Braden, "Tussles in Cyberspace: Defining Tomorrow's Internet", Sigcomm 2002, extended version, Transactions on Networking, v.33, no. 5.
- [7] M. Demmer, Fall, K., Koponen, T., and Shenker, S. Towards a Modern Communications API, Proc. HotNets VI, 2007.
- [8] DTMF, "Common Information Model (CIM) Specification" V. 2.2, 1999, DSP0004.
- [9] Enterprise Management Associates, Practical Autonomic Computing: Roadmap to Self Managing Technology, White paper prepared for IBM, January 2006.
- [10] Eugster, P. T., Felber, P., Guerraoui, R., and Kermarrec, A. The Many Faces of Publish/Subscribe, ACM Computing Surveys, 35(1), 2005, pp. 114-131.
- [11] N. Goodman, V. Mausingshka, D. Roy, K. Bonawitz, J. Tenenbaum, "Church: a language for generative models", Uncertainty in Artificial Intelligence 2008.
- [12] Greenberg, A., Hjalmytsson, G., Maltz, D. A., Myers, A., and Rexford, J. A Clean Slate 4D Approach to Network Control and Management, ACM SIGCOMM Computer Communications Review, 35(5), 2005.
- [13] M. Huang, A. Bavier, L. Peterson, "PlanetFlow: Maintaining Accountability for Network Services", Operating Systems Review, Vol 40, No 1. January 2006
- [14] Y. Huang, N. Feamster, A. Lakhina, and J. Xu, "Diagnosing Network Disruptions with Network-Wide Analysis", ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, 2007.
- [15] IBM, An architectural blueprint for autonomic computing, White paper, fourth edition, June 2006.
- [16] V. Jacobson, "A new way to look at networking", at <http://video.google.com/videoplay?docid=-6972678839686672840>
- [17] P. Jokela, A. Zahemszky, S. Arianfar, P. Nikander, C. Esteve, "LIPDIN: Line Speed Publish/Subscribe Inter-Networking", PSIRP Project, TR09-0001, Jan. 2009.
- [18] A. Keller, Towards Autonomic Networking Middleware, Slides from Keynote talk at International Workshop on Next Generation Networking Middleware, May, 2005.
- [19] J. Kephart and D. Chess, *The Vision of Autonomic Computing*, IEEE Computer, (), January, 2003
- [20] J. Kopena and B. Loo, "OntoNet: Scalable Knowledge-Based Networking", 4th Workshop on Networking Meets Databases, 2008.
- [21] T. Koponen et al., "A Data-Oriented Network Architecture", ACM SIGCOMM, 2007
- [22] G. Lee, "CAPRI: a common architecture for distributed probabilistic Internet fault diagnosis", MIT Department of Electrical Engineering and Computer Science, May 2007.
- [23] J. Li, "Agent Organization and Request Propagation in the Knowledge Plane", MIT Department of Electrical Engineering and Computer Science, May 2008.
- [24] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy and A. Venkataramani, "iPlane: An Information Plane for Distributed Systems", OSDI 2006.
- [25] Network Science and Engineering Council, "Network Science and Engineering Research Agenda", Release V. 1, July, 2009.
- [26] K. Park, V. Pai, "CoMon: A Mostly-Scalable Monitoring System for PlanetLab", Operating Systems Review 40(1): 65-74 (2006)
- [27] L. Peterson, A. Bavier, M. Fluczynski, S. Muir, "Experiences Building PlanetLab", OSDI '06, 2006.
- [28] A. Phillips, "The SPiM Language", <http://research.microsoft.com/en-us/projects/spim/>, 2007 (found Aug. 2009).
- [29] K. Sollins, "Designing for Scale and Differentiation". ACM SIGCOMM FDNA 2003 Workshop, Karlsruhe, August 2003.
- [30] K. Sollins, J. Van Dyke, "Linking in Global Information Architecture", 4th WWW Conference, Dec. 1995.
- [31] A. Tripathi, "Ajanta (Mobile Agents Research) project", <http://research.microsoft.com/en-us/projects/spim/> as of Aug. 2009.
- [32] D. Trossen (ed), "Architecture Definition, Components Descriptions and Requirements", at <http://www.psirp.org>, 2009.
- [33] W3C, "OWL Web Ontology Language Reference", W3C Recommendation, Feb. 10, 2004.
- [34] Wawrzoniak, M., Peterson, L., and Roscoe, T. Sophia: An Information Plane for Networked Systems, (HOTNETS-II) Proc. 2nd Workshop on Hot Topics in Networks, 2004, pp. 15-20.