Research Challenges in Future Networks: A Report from US-Japan Workshop on Future Networks

Deep Medhi
Department of Computer Science & Electrical Engineering
University of Missouri–Kansas City
dmedhi@umkc.edu

Peter A. Freeman
College of Computing
Georgia Institute of Technology
peter.freeman@mindspring.com

This article is an editorial note submitted to CCR. It has NOT been peer reviewed. The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

A US-Japan Workshop on Future Networks was held in Palo Alto, CA on October 31 - November 1, 2008. This workshop brought together leading US and Japanese network researchers and network research infrastructure developers who are interested in future networks. The focus was on research issues and experimental infrastructure to support research on future generation networks. The goal was to foster cooperation and communication between peers in the two countries. Through this workshop, a number of research challenges were identified. The workshop also made recommendations to: create a new funding mechanism to foster special collaborations in networking and experimentation, extend current national testbeds with international connectivity; and urge the respective governments to exert strong leadership to ensure that collaborative research for creating future networks is carried out.

Categories and Subject Descriptors

A.m [**General**]: Miscellaneous—Workshops

Keywords

Future Networks; Research Challenges; Collaborative Research.

1. INTRODUCTION

A US-Japan Workshop on Future Networks, sponsored by the National Science Foundation (NSF) of US, was held in Palo Alto, CA on October 31 - November 1, 2008. This workshop brought together leading US and Japanese network researchers and network research infrastructure developers who are interested in future networks. In Japan, future network initiatives include AKARI, an Architecture for New Generation Network, research projects funded by National Institute of Information and Communications Technology (NICT) of Japan for new generation network technologies, and Japan Gigabit Network (JGN2Plus), to support fundamental research and experimental in networking. In the US, the main initiatives include NetSE (Network Science & Engineering), FIND (Future Internet Design), and GENI (Global Environment for Network Innovations) programs, supported by the National Science Foundation.

The Workshop built on the outcomes of an earlier workshop held in Tokyo in January 2008 and provided a forum for peer researchers from the US and from Japan to explore different possibilities for cooperation in network design and engineering. The focus was on discussing research issues and experimental infrastructure related to future generation networks. The goal was to foster cooperation and communication between peer researchers and research infrastructure developers in the two countries.

The two-day workshop started with brief welcome remarks by Ty Znati, Division Director, Computer and Network System (CNS) Division, NSF and Hiroshi Kumagai, Executive Director, New Generation Network Research Center, NICT. They both emphasized the importance of this workshop to identify research challenges and potential collaborations between researchers in the two countries. This was followed by opening remarks by Peter Freeman and Tomonori Aoyama, General co-chairs of the workshop. They emphasized the importance and the necessity of collaborations that are increasingly necessary to solving deep scientific problems, and that often collaborations among individuals in different fields are critical. Future networking research is of this type, involving researchers from both countries who must come together. General co-chair Aoyama further elaborated on recent activities in Japan on new generation networks as well government policy, interests, and funding priorities for future networking. For NSF, Suzi Iacono then elaborated on NSF funding mechanisms and opportunities for international collaboration, while pointing out the National Science Board's recent report that emphasized the need for a more effective NSF role in international science and engineering.

The above presentations were followed by twelve single-track presentations by leading researchers from the US and Japan. These presentations encompassed future network testbeds and infrastructures, future networking issues and requirements, and the potential for future collaborations. On the morning of the second day, four breakout sessions, each led by two co-chairs (one each from the US and Japan), were held; participants were assigned to sessions. The session co-chairs then reported their findings to the entire group for further discussion and feedback, followed by identification and summarization of research issues and potential future collaborations and goals. The agenda, the presentations, and the list of participants can be found at [1].

As a special note, this workshop was dedicated to the memory of Masaki Hirabaru, a leading networking researcher from Japan, who died in 2008.

2. RESEARCH CHALLENGES

To identify research challenges, breakout sessions were organized around four themes: 1) network fundamentals, 2) network architecture and design, 3) socio-economic and environment-aware network services, and 4) experimental infrastructure for future network design. In the following, we elaborate on these themes, their focus, and the discussion outcomes.

2.1 Network Fundamentals

Designing networked systems and services often involves making a number of assumptions and using mathematical models to understand the behavior of these systems and their performance in a variety of settings. Emergence, defined as the formation of complex but regular patterns from the interaction of the many simple parts of a system, is a key property of complex network systems. It is this element of regularity in the emergent behavior that distinguishes complex network systems from complicated and chaotic systems. Emergent behavior of complex network systems, however, cannot be predicted merely by understanding the behavior of their individual elements, or from understanding the interactions between these elements. This breakout session focused on models, tools and mechanisms to advance our understanding of key properties of complex network systems and their behavior.

- Are there “laws” that govern the structure and consequently the behavior of complex networked systems?
- Can a theory be developed to assess the vulnerabilities and fragilities inherent in complex networked systems so that we can design them to have predictable or adaptable behavior?
- Biological systems are often able to efficiently adapt to environmental changes as they occur and often succeed in carrying out their tasks efficiently and reliably, even in the presence of multiple component failures. Would network models, inspired by biology and other systems, lead to reliable, evolvable, self-manageable and self-healing networks?

As a result of the discussion in this breakout session, several common research challenges were identified as listed below:

1. How do we identify methodologies that apply to future network design for which inspiration could come from such areas as electromagnetics, biology, economics, game theory, statistical physics, or traditional methods like queueing theory, heavy traffic models, or optimization theory.
2. How do we create metrics of “goodness” for models and identify domains where they apply? Which models give insight into different aspects of network performance?
3. How do we create environments to enable architecture design and performance modeling at the same time?
4. How do we develop synergistic models that contain complex dynamic theory, optimization theory, cognitive or machine learning theory? Can new models produce satisficing results?
5. Is there a relationship between power law topologies and power law in traffic?
6. How do we create a common language to describe inputs, outputs, and metrics for large complex networks?
7. How do we obtain a specific desired network behavior in self-organizing systems?

The next issues addressed were the impediments felt by the researchers of both countries to conducting the above research and how they could be overcome.

The main impediments identified were: 1) how to validate models, 2) how to prune and reduce the number of modeling approaches, 3) how to decompose large complex systems into more understandable components, 4) how to identify proper objectives and metrics, 4) how to ensure that modeling methodologies guide the design of future networks.

Overcoming these impediments is itself challenging as well. In spite of this, this session identified that the main key to overcoming these impediments were 1) the need for futuristic experimental infrastructure, 2) the need for forming research teams with different and complementary expertise, and 3) the need for gathering information on dynamics from current networks that may in turn enable an understanding of the future impediments.

2.2 Network Architecture Design

The design and deployment of future networks requires rethinking current network design principles, exploring new paradigms that go beyond current circuit- and packet-switching techniques and the development of strategies to address a range of challenges and opportunities, such as security and privacy, mobility, and ease of management and operation. The focus of this breakout session was on architecture and design principles, protocols and mechanisms for more reliable, available, secure and evolvable future networks and services. The questions discussed included:

- Is the traditional concept of layering fundamental to network design, and if so, what layering-based framework and approaches will enable cross-layer optimization while preserving efficiency, modularity and ease of design and management?
- What network design principles and abstractions are required to support functions such as any-to-any information routing and dissemination, addressing, location management and identity management in heterogeneous, data-intensive, wireless and mobile networking environments? How do we conceive these network services to be secure and privacy-preserving?
- Is the end-to-end argument still applicable? If not, what minimal set of application-specific functionalities can be implemented in the network core to meet the demands and requirements of emerging and future network applications?
- What impact will current and emerging technologies, such as cognitive wireless devices and programmable

optical links, have on network architecture, protocols and services?

- How do we transition to new architectures, especially if they are "clean slate"?

As a result of the discussion in this breakout session, three broad research challenges were identified as listed below:

1. How do we create networks for the masses - making infrastructure to support society broadly?

2. How do we effectively organize network components that allow dynamism and virtualization?

3. Is there a science of network architecture, and what are its properties?

The group then considered impediments and how to overcome them.

The major impediments identified were: 1) understanding and having predictable requirements, 2) addressing conflicting (and sometimes directly opposing) requirements and goals, 3) system and design complexity, and 4) deployment hurdles. To overcome these impediments, a multitude of dimensions must be addressed in both countries. These include having a vision of a future data plane that is extensible and that will accommodate multiple use technologies, developing specification languages of sufficient generality, and having applications that will drive the transition to newer architectures.

2.3 Socio-economic and Environment-Aware Network Services

Future networks must consider socio-economic values and requirements, and must be resource-efficient and environmentally friendly. This breakout session focused on the design principles, models and tools to enable the design, deployment, management and operation of future networks that meet these goals. The questions discussed included:

- How would social, legal and economic requirements shape or impact the overall design of future network architectures to support network functions and services in an open and free society?

- How can network protocols, network services, traffic management and data control be designed to be environment friendly? Can we define actionable goals and articulate strategies, with clear understanding of roles, deliverables, priorities and capabilities, which can lead to significant reduction of power consumption and minimization of undesirable chemical element emission?

- What models can be used to advance our understanding of the relationship between network processing and power consumption? What models can be used to assess total energy consumption and understand the impact of optical and wireless technologies on green internetworking?

This breakout session identified several research challenges as listed below:

1. How do we consider end-to-end energy profiling that can be useful in future networks?

2. How do we design power-aware network architecture that is "not-always-on" without impacting the lifetime issue of devices?

3. How does dynamic power-rate adaptation impact quality of service?

4. How does application usage profile for different socio-economic strata of the society impact power consumption on end-devices and what problems need to be addressed to optimize power usage?

5. How do we use fewer devices with optimal power usage while meeting service requirements?

6. How do we take an interdisciplinary approach to understanding various networking requirements?

7. Can we forecast power needs from forecasted network traffic?

8. How can we ensure that RF pollution is contained?

9. How do we make systems that are safe inside of ecosystems?

There are several impediments to overcoming these challenges. It is important to understand power consumption and usage behavior in order to do energy profiling. Secondly, it remains a challenge to do realistic measurements for the purpose of energy-aware goals. Thirdly, any measurement or assessment significantly differs from one region of the world to another as well as due to socio-economic differences.

Overcoming these impediments requires concerted efforts in understanding socio-economic and geo-political differences so that environment-friendly networking can be properly addressed.

2.4 Experimental Infrastructure for Future Network Design

Experimental infrastructure is crucial to supporting and advancing our understanding of how network and distributed services are designed to meet societal values, network policy and economics. The objective of this breakout session was to provide leading networking experts with an opportunity to draw and confirm common views about experimentation and testbeds for future networks. The questions addressed included:

- What challenges must be addressed to enable the design, deployment, administration, operation and management of global experimental infrastructure to support collaborative network research, in a real-world environment unconstrained by current testbed limitations or existing network regulations and policies?

- Testbeds for future networks must enable the design, experimentation and prototyping at scale of advanced network protocols and architectures, and control and management paradigms, including testing of the interplay between the socio-technical, economic and policy issues related to these networks. Are virtualization, programmability and federation across all layers and at scale sufficient to achieve this goal? If not, what other challenges must be addressed to enable collaborative experimentation at scale in a secure and privacy-preserving manner?

There are major challenges in building an experimental infrastructure for future networks. For instance, such an infrastructure will require the deployment of equipment from vendors, raising the challenging issue of not having access to open software codes - an important ingredient for building a flexible infrastructure - and it must allow virtualization and programmability. This brings up a number of issues:

1. How do we accomplish virtualization and heterogeneity at a global scale, with diverse collection of resources?
2. How do we extend virtualization to wireless edges?
3. How do we implement virtualization across layers of a protocol stack?
4. How do we advertise, discover, and manage resources in a virtualized environment?
5. How do we share resources in a fair way?
6. How do we map applications' needs and services in a virtualized environment?
7. How do we enforce security and access control in a scalable manner?
8. How does virtualization impact business models and service providers, and how to price services?
9. How do we define a testbed that supports yet-undefined experiments without interfering with them?

To understand these issues, there is a critical need for a federated experimental platform. Furthermore, the funding mechanism for such an experimental platform and the exact role of government agencies and the researchers are not yet clear, and how different countries can become involved is left open.

Despite all the hurdles, progress must be made, perhaps with small, concrete steps. Incentives must be provided to allow service providers and vendors to be part of the experimental facility. Where possible, simulations may be carried out to understand parts of the facility providing a complementary role. Open communication among all stakeholders, led by researchers, can help overcome these hurdles.

3. COMMON MAJOR THEMES AND RECOMMENDATIONS

From the individual presentations and breakout sessions, some common themes came to the surface. First and foremost is that researchers are very interested in envisioning future networks and the capabilities that would be needed so as to serve broader socio-cultural communication needs. A general belief was expressed that a robust networking framework is important to help ensure and enable future networking and services that are flexible, yet robust.

The attendees were also of the opinion that a strong collaboration between researchers in the US and Japan will be important in formulating the requirement and the design for future networks. As a part of the breakout sessions, potentials for collaborations were discussed and pursued. It is clear that creating the right atmosphere for scholarly exchanges is critical. In addition, the current national testbeds must be extended to bring in the international connectivity that can help foster new international experimentation. The researchers also felt that early starters for international experimentation may build on Planetlab and Openflow, while scholarly collaborations must be fostered over the next twelve months in areas such as virtualization, policy, control and data plane interfaces, security and service abstractions. Through such collaborations, common methodologies (such as for network usage or security policies) can be developed.

It was also commonly believed that the governments and their research funding agencies play a critical role. More importantly, having clear funding mechanisms for international collaboration and one that provides substantial funding (i.e., not just travel funds or support for one or two students) are extremely important. In fact, the research community urges funding agencies in both countries to consider this issue proactively and to give it the highest importance by issuing special solicitations for research and infrastructure for international collaborations. This will greatly help foster international collaborations quickly and help jumpstart new research directions in future networking. In this regard, an NSF PIRE (Partnership for International Research and Education) type program for international collaborations in future networking is strongly desirable.

To summarize, the workshop made the following primary recommendations:

- Create a new funding mechanism to foster special collaborations in networking and experimentation. While the NSF Partnership in International Research & Education (PIRE) program exists, a new funding mechanism for fostering collaborative research in future networking and experimentation at the international scale is critically needed due to the economic importance of future networking.
- Extend current national testbeds with international connectivity to help foster multi-nation collaborative experimentations in future networks; for example, federate between JGN2Plus and GENI.
- Governments in both countries should continue to exert strong leadership by ensuring that collaborative research and support for global environment for large scale experimentation for creating future networks are being carried out.

4. SUMMARY

The co-chairs and organizers felt that overall, the workshop was a success and that the two days generated many productive discussions. The Japanese have offered to host a follow-up workshop in Japan, and the detailed plan is to be discussed and determined by NSF and NICT. In the meantime, strong efforts should be made by NSF, NICT, and each participant to develop more and stronger collaborations in the area of networking research infrastructure. Current collaborations and future candidates were discussed at the workshop; the goal is to have at least six active collaborations by the time of the next workshop. A future workshop should focus on the results of infrastructure work, not just plans, and should primarily be about truly new approaches that address ideas outside the box of current networking fundamentals.

5. ACKNOWLEDGEMENT

This Workshop was supported by Grant No. 0853414 (Deep Medhi, PI) from the Computer & Network Systems (CNS) Division of the National Science Foundation. Travel expenses for invited US participants were covered through this grant. Some of the invited workshop participants from Japan were reimbursed for travel expenses by NICT.

Peter Freeman and Tomonori Aoyama served as the general co-chairs of the workshop while Deep Medhi and Hiroaki Harai served as the workshop arrangement co-chairs.

The workshop web-site [1] was hosted by the Department of Computer Science & Electrical Engineering, University of Missouri–Kansas City.

A Wiki page [2] was created by James Sterbenz, and is hosted by the Information and Telecommunication Technology Center, The University of Kansas.

This report is based on the final report submitted to NSF; the final report can be found at [1]. The authors thank Tomonori Aoyama, Jim Kurose, and Joe Touch for carefully reading a draft of the final report and providing valuable feedbacks.

6. REFERENCES

[1] US-Japan Workshop on Future Networks, web-site: http://www.csee.umkc.edu/us-japan-fnw/

[2] Wiki page for the US-Japan Future Networks Workshop (restricted access) http://wiki.ittc.ku.edu/nsf-us-japan_wiki/