The New Generation Network (NwGN) Project:
Its Promises and Challenges

Keynote Speech presented at the NICT New Vision Symposium

November 9, 2011 at Tokyo Conference Center-Shinagawa

Hisashi Kobayashi

The Sherman Fairchild University Professor Emeritus of
Electrical Engineering and Computer Science,
Princeton University, Princeton, New Jersey, USA
and
Executive Advisor
National Institute of Communications and Information and Technology

It is a great honor to be invited to this symposium. I thank President Miyahara, Vice President Enami and the organizing committee for providing me this opportunity. I am delighted that this symposium, which was originally scheduled in the past June is finally held today.

Slide 2: Outline of the presentation

Here is the outline of my talk

I. What is NwGN and Why?
II. What made the Internet successful?
III. An end to the End-to-End design argument?
IV. Network Virtualization
V. AKARI Architecture and JGN-X
VI. FIA and Testbed efforts in the U.S. and elsewhere
VII. What should be Japan's strategy from now on?
VIII. What should we learn from Mr. Steve Jobs?

Slide 3: What is NwGN ?

The NwGN project is a flagship project, so to speak, of the networking research in Japan.. The NwGN intends to make a revolutionary jump from the current Internet. Its purpose is to design a new architecture, and implement and verify it on a test bed so as to enter its experimental phase around 2015. Such a fresh approach is often referred to as a "clean slate" approach.

Slide 4: Why NwGN?

The current Internet has been constantly and successfully transforming itself to support new applications and accommodate ever-increasing new users, a great portion of which is now mobile terminals such as smart phones, laptops, sensors and the like [1].

The speed, at which the Internet has evolved and expanded in its usage and applications over the past two decades since the commercialization of the Internet launched in 1988, is unprecedented in the history of human civilization. Forty years ago the original designers of the Internet (called ARPANET then) assumed that network terminals were fixed and all users were trustworthy. Consider, however, the explosively growing network traffic today, the increasing difficulty in protecting the Internet against sophisticated cyber attacks, and the trend that a majority of terminal devices is now mobile as is the case with smart-phones, laptop PCs and sensors. Then, it should be rather obvious that the NGN (Next Generation Network)—which is merely an extension of today's Internet, designed to provide triple-play services, i.e., (1) telephony, (2) internet, and (3) television, with its world wide standards already set up and now has already entered the phase—will hit its limit sooner or later.

The NwGN project will aim at a revolutionary change so as to meet societal needs of the future. AKARI is the architecture of such a network and JGN-X is a testbed, on which we will implement and verify the new architecture.

Slides 5 & 6: Requirements of NwGN

There are numerous requirements that we need to take into account concerning network services of the future: Here is a list of what I consider as requirements for the NwGN.

1. **Scalability** (users, things, data, traffic**)**
2. **Heterogeneity** and **diversity**
3. **Reliability** and **resilience**
4. **Security protection**
5. **Mobility management**
6. **Performance**
7. **Energy** and **Environment**
8. **Societal** needs
9. **Compatibility (with today's Internet)**
10. **Extensibility (for the unforeseen and unexpected)**

The NwGN's "Vision statements" address all of the above issues in varying degrees[2]. I will discuss in this talk some highlights of **AKARI** (the future Internet architecture that the NwGN group has been pursuing with Dr. Harai as its leader) and **JGN-X** (a testbed that Prof. Shimojo et *al*. and his collaborators in the world have been developing as experimental facilities to test new architecture, protocols and applications). But I should first clearly define, for the general audience, such terms as "architecture" and "protocol," and then briefly review what has made today's Internet so successful, and which features of the current Internet should be retained and which should be discarded.

I. What are Architectures and Protocols ?

A **network architecture** is a logical and structural framework of a network, and specifies the network's components and their functional organization and configuration, its operational principles and procedures. The architecture also specifies data format used in its operation. An architecture is different from its *implementation* in that it is not tied down to a specific technology to be used. In other words, an architecture should be, in principle, independent of technologies.

A network architecture is specified by **protocols**. In diplomacy, the term "protocol" refers to the set of rules, procedures, conventions and ceremonies which relate to the relations between nations. In computer communications, this term is used to mean a set of rules that enable communications between two computers/end devices.

The architecture of the current Internet has been predominantly specified by **Internet Protocol** (**IP**), which essentially describes how packets are routed within an internetwork. Network services provided by IP alone are said to be **connectionless** in that it does not provide orderly transmission of packets between a given pair of end points. In order to provide **connection-oriented** services, we need to specify another protocol that the end points must adhere to, and such a protocol in the Internet is called **Transmission Control Protocol** (**TCP**) described in a seminal paper [2] published by Cerf and Kahn in 1974. TCP runs on top of IP and this **TCP/IP** protocol **suite** has been serving as the foundation protocols. Other commonly used protocols include, besides TCP/IP, **File Transfer Protocol** (**FTP**), and **Hypertext Transfer Protocol** (**HTTP**). HTTP is the set of rules for transferring files (text, graphic images, sound, video, and other multimedia files) on the World Wide Web. HTTP is an application protocol that runs on top of the TCP/IP suite of protocols.

II. What made the Internet successful?

It is remarkable that the TCP/IP protocol suite, a nearly four decade old architectural concept, remains still in effect, despite the profound changes in the speed, capacity and cost that various network component technologies have undergone in the past four decades [1]. In this revolutionary digital age, forty years are a long time. So what has made the Internet continuously grow and thrive for so long? Slide 8 illustrates a simplified view of the conventional **telephone network**, in which all intelligence resides within the core network. The end devices are dumb telephones

Slide 9 illustrates a **TCP/IP** based Internet. The design principle of TCP/IP is often referred to as the "**End-to-end** (**E2E**) **design**" principle. Roughly speaking, this principle says that the core network should be a *dumb* network, and end-points of the network should provide all functions/services that are required to support any application.

The most important implication of this E2E design principle is the architecture's complete **openness,** that is, anyone could develop a new application that runs on the IP network because the dumb core network is readily accessible and easy to deal with. Such accessibility is not available in the telephone network, whose highly centralized control structure does not allow end users to access the interior network. The open nature of the Internet architecture is the main reason why the Internet won the race to multimedia services against the concerted world wide effort of the group of telecommunication carriers that were developing **B-ISDN** (broadband integrated services data network) with ATM (asynchronous transfer mode) fast packet switching. As many of you may recall, B-ISDN was once the grand vision of multimedia services of the 21st century outlined by the international standardization committees of telecommunications carriers in the 1980s-1990s.

III. An end to the End-to-End design argument?

If we review the above list of ten requirements for future network services, the E2E design approach cannot adequately cope with many of these requirements. The E2E design would make the core network intrinsically inefficient (i.e., wastes network resources such as bandwidths), often slow (i.e., invokes unnecessary retransmissions in case of errors that may occur in noisy links, such as radio links) and internally insecure (because a dumb network cannot protect itself from attacks).

Therefore, we should no longer adhere to the E2E. In the interest of time, I will not dwell on this further. Please refer to my slides and the extended abstract presented in my keynote entitled "An End to the End-to-End Arguments," which was presented at a Euroview conference held in Wurzburg, Germany in July 2009 [4].

IV. Network Virtualization

If we must do away with the E-2-E design approach, which has greatly helped the Internet enhance itself and has allowed fast development of new applications, how can we provide a similarly excellent environment for enhancements and applications development in the future Internet?

The most important recent development in the networking research, especially that of its experimental facilities or "testbeds," is "**network virtualization**." The original notion of "virtualization" in computer technologies goes back to circa 1960, when *virtual memory* was introduced in Atlas machine of the University of Manchester, UK. In 1972, IBM introduced VM/370, a *virtual machine* (VM) operating system that ran on System/370. In the last decade, IT (information technology) departments of enterprises have begun to adopt a variety of virtualization technologies available as commercial products, ranging from *server virtualization, storage virtualization, client (or desktop) virtualization* to *software virtualization* (e.g., allowing Linux to run as a *guest* on top of a PC that is natively running a Microsoft Windows operating system). Such virtualization techniques allow multiple users and applications to dynamically share physical resources. Thus, they increase resource utilization and/or reduce electric energy consumption, as well as simplifying complex administrative operations of IT.

Network virtualization chooses a subset of a collection of *real* (or *physica*l) *resources* (routers, end users, links,etc.) and *functionalities* (routing, switching, transport) of a real network (or multiple real networks) and combines them to form a **logical network** called a **virtual network**. Network virtualization allows multiple instances of network architectures and/or protocols to run on separate virtual networks, facilitating the design, implementation and validation of a novel network architecture and protocols in a realistic network environment, with minimal interference to existing network services. This approach is currently actively pursued by a number of research initiatives on novel architectural designs of future networks and implementation of experimental testbeds. Virtual networks take different forms, depending on specific layers to which virtualization is applied. Prof. Akihiro Nakao of the University of

Tokyo, who also participates in the NwGN and JGN-X projects, is a leading expert on this topic. Please refer to his papers available on his website and an excellent tutorial article on his "Virtual Node Project" that appeared in NICT News' June 2010 issue [6].

V. AKARI Network Architecture

The AKARI network architecture project was launched in 2006 as a part the NwGN effort, with participation by a dozen professors from Japanese universities who are active in networking research. If you are interested in details of the accomplishment of the AKARI efforts, you can find the conceptual documents of AKARI architecture, published both in English and Japanese [2].

VI. JGN-X Testbed

NICT's test bed effort for NwGN is called JGN-X, which is an evolutionary outgrowth of JGN (Japan Gigabit Network) that started in Year 2000 as a test bed for large capacity. As its speed and capacity got increased, the name changed to JGN2, JGN2 plus, and now JGN-X, where X stands for "eXtreme." Figure 3 shows the Roadmap of JGN-X together with NwGN.

VII. FIA and Testbed efforts in the U.S. and elsewhere

It goes without saying that comparable efforts are going on in other parts of the world. In the United States, the Networking Technology and Science (NeTS) program of the National Science Foundation (NSF) initiated the **FIND** (Future Internet Network Design) initiative in 2005 as a long-term networking program and has funded about 50 research projects on design aspects of the future Internet.

As its testbed efforts, the NSF-sponsored **GENI** (Global Environment for Network Innovations) program (2005-present) states its mission as "a unique virtual laboratory for at-scale networking experimentation where the brightest minds unite to envision and create new possibilities of future internets."

GENI is managed by Mr. Chip Elliot of BBN Technologies, who holds quarterly meetings/workshops, called GEC (GENI Engineering Conference), where the researchers and students of each group funded by the GENI office are required to report their progress of the past three months. I have attended four GEC meetings in the past three years, and I have been impressed by how fast each of the four testbed groups (called "GENI Control Framework" or simply "clusters") has been making progress. The following four clusters (lead institutions) are currently supported: **PlanetLab** (Princeton University), **ProtoGENI** (Univ. of Utah), **ORCA** (Duke University and RENCI-Renaissance Computing Institute) and **ORBIT** (Rutgers University).

Furthermore, In 2010, NSF launched the **FIA** (Future Internet Architecture) program, which supports fewer but larger team projects than the many small projects funded by FIND. Currently the program support four projects: **MobilityFirst** (Rutgers and 7 other universities), **Named Data Networking** (**NDN**; UCLA and 10 other universities), **eXpressive Internet Architecture** (**XIA**:CMU and 2 other universities), and **NEBULA** (U. of Penn and 11 universities). Each FIA program has its own comprehensive website where you can find more information than you possibly digest. A recent survey paper in the July 2011 issue of IEEE Communications Magazine provides a good introduction to the FIA , GENI and EU's programs. The article also allocates about a half page to AKARI and JGN-X.

Most recently, the NSF and President Obama's office launched what is called the **U.S. Ignite Initiative**. It is my understanding that this initiative is a response by President Obama to the challenges from other countries, where the broadband Internet connections to homes are much faster than in the U.S. (see e.g. [8]) and private-public partnership (PPP) programs that have been recently placed in EU (see the paragraph below).

In Europe, a collaboration of **FP7** (the Seventh Framework Programme) on Future Internet research is referred to, somewhat confusingly, as **Future Internet Assembly** (**FIA**). The **EIFFEL** (European Internet Future for European Leadership) support action was launched in 2006 "to mobilize European researchers to discuss and debate on the future of the Internet towards the development of the future networked society." In May 2011, the **Future Internet Private-Public Partnership** (FI-PPP) was launched by the European Commission as a fourth PPI initiative, as part of the European Economic Recovery Plan. (The three previous PPIs are "PPP1: Factory of the Future," "PPP2: Energy-efficient Buildings," and "PPP3: Green Cars.")

Germany has been sponsoring the **G-Lab** (German Laboratory) through BFBM (Bundesministrerium für Bildung und Forschung; Federal Ministry of Education and Research) in addition to their participation in the aforementioned EU efforts

South Korea, which enjoys the fastest Internet connections in the world already, and intends to connect every home in the country at one gigabits/sec by the end of 2012, has also the Future Internet program in place.

VIII. How should Japan cooperate and compete in the Future Internet era?

So where does Japan's effort on the future Internet and related ICT (information and communications technology) segments stand vis-à-vis similar efforts in the U.S., Europe, South Korea and others? It is a meaningful question to pose because the ICT sector will continue to be a significant part of the overall economy of a nation, its international competitiveness, and wellbeing of its citizens. This sector includes other rapidly expanding technologies such as **cloud computing** and futuristic applications such as **smart city** (and the smarter city). Counting the number of professors, researchers and graduate students engaged in networking research may serve as one indicator. Comparison of R&D budgets allocated to funding programs in these areas may be another. I don't have any of these statistics, but it seems rather clear that the networking community dwarfs the U.S. by a wide margin, perhaps by an order of magnitude? The U.S. networking community can afford pursuing four or more different architectures in parallel, as seen in the FIA programs and other project supported by the NeSE program, whereas in Japan there seems no other internationally visible effort than NICT's NwGN and AKARI efforts. Therefore, it is absolutely critical that we keep abreast of the FIA efforts in the U.S. and elsewhere. At the same time we have to ensure that our peers in other countries are well informed of what we have been doing so that some of our key ideas and achievements should be incorporated in the world wide standard of the FIA. It is encouraging that Prof. Nakao and Dr. Kawai made presentations on the NwGN project, (AKARI overview, network virtualization and JGN-X) at a plenary session of the GEC 11 meeting in July.

It is not so easy to substantially increase the size of networking community and/ or to greatly improve the quality of people in a short term, unless we should come up with a much larger budget that would allow us to recruit a number of top-notch researchers and project leaders from abroad, which Japan should have done a few decades ago when Japan was in a financially stronger position. There are a few things, however, that the Japanese research community can do with little additional resources. First, we should make it mandatory to prepare **presentation slides in English** even if all participants are expected to understand Japanese. Preparing your slides and papers in Japanese is a sure way to exclude non-Japanese from the meeting or conference. Second, we should give a lot more opportunities for **graduate students** and **post-doc fellows** to stand up on a big stage like this and report their exciting projects. Thirdly, we should encourage **our graduate students and young researchers to spend a summer, a semester or a year at our collaborating institutions abroad**. A common mistake that Japanese professors or managers often make in such an arrangement, however, is to let our students or research staff to continue working on the same topics currently pursued here. That would kill the whole purpose of exchange programs. Let them instead partake the project of their host institutions and fully engage in the projects of the host institution. Reciprocal efforts are also important. Invite more graduate students and post docs from abroad to NICT and Japanese universities. Exposing our young people to people with different background and alternative ways of thinking is the most effective, and perhaps the only way to generate innovative ideas by breaking away from a monolithic environment that the Japanese institution tends to create.

IX. What can we learn from Mr. Steve Jobs?

Two months ago the world lost Mr. Steve Jobs, an ingenious product designer, a forceful manager and charismatic CEO and a great deal has been talked about his enormous achievements, Apple's uncertain future, etc. So it may be fitting for me to add a brief but inspiring anecdote to what you may have already learned about him.

I had an opportunity of meeting with Mr. Steve Jobs around 1987 or 88 (so he was quite young, 32 or 33 years old), when he, as the CEO of his new company NeXT, was developing NeXTcube, a workstation for college students and business people. He visited several university campuses in the country, including Princeton, to promote his forthcoming product. He was also genuinely interested in getting feedback from Princeton's professors and students after he made a presentation about his new product in a fully packed auditorium.

As the then Dean of Engineering , I had the pleasure of hosting a dinner for him at Prospect House, Princeton's Faculty Club. As soon as we were seated at the table, and waiters started serving soups and appetizers, he stood up saying, "Excuse me. I have to make a phone call," and walked to a telephone booth outside the dining room (it was well before a cellular phone came into existence). After ten minutes or so, when the rest of us already started eating, Steve returned to the table, and told me that he talked to his engineers about some new idea he got at Princeton that afternoon, telling them to make some design changes immediately. This episode is consistent with what I later learned about him. He was not only a brilliant architect of a grand concept of a new product, but also was engaged in every detail of the product design and its marketing plan. He did not want to waste his engineers' time by postponing his phone call until the dinner was over or until the following day when he would return to his office.

I am sure that many of you have seen on You Tube Steve Jobs' famous speech at Stanford University's Commencement in 2005. If you haven't, please do so by all means. It is the most moving and inspiring speech I have ever listened to. Actually I viewed the video a few times. Just type in Google "Steve Jobs' speech at Stanford's Commencement." There is a You Tube video with a Japanese subtitle or Japanese text [6]. His speech at Stanford is definitely a must for everyone here.

References

[1] H. Kobayashi, "Forword," in D. Raychaudhuri and M. Gerla (edtitors), *Emerging Wireless Technologies and the Future Mobile Internet*, Cambridge University Press, 2011.

[2] New Generation Network Architecture AKARI: Its Concept and Design (ver2.0) September, 2009.

[3] V. G. Cerf and R. E. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Trans. on Comm.*, 22(5), pp. 637-648, May 1974.

[4] J. H. Saltzer, D. P. Reed and D. D. Clark, "End-to-End Argument in System Design," *ACM TOCS*, Vol. 2, No. 4, November 1984, pp. 278-288

[5] H. Kobayashi, "An End to the End-to-End Arguments," *Euroview 2009*, Wurzburg, Germany, July 2009

[6] A. Nakao, "Virtual Node Project: Virtualization Technology for Building New-Generation Networks," *NICT News*, June 2010, No. 393, June 2010, pp. 1-6.

[7] J. Pan, S. Paul and R. Jain, "A Survey of the Research on Future Internet Architectures," *IEEE Communications Magazine*, July 2011, pp. 26-35.

[8] New York Times "Home Internet May Get Even Faster in South Korea," February 21, 2011.
http://www.nytimes.com/2011/02/22/technology/22iht-broadband22.html?_r=3

[9] Steve Jobs's Speech at Stanford University Commencement, 2005

http://www.youtube.com/watch?v=D1R-jKKp3NA

http://sago.livedoor.biz/archives/50251034.html