

New-Generation Network Architecture: Its Opportunities and Challenges

A Keynote Address

The Opening Ceremony of
Keihanna Research Laboratories
The National Institute of Information and Communications Technology (NICT)

June 23, 2008

Hisashi Kobayashi
The Sherman Fairchild University Professor of
Electrical Engineering and Computer Science
Princeton University, Princeton, NJ, USA

Abstract: After observing that there is an analogy between the NICT's new-generation network or AKARI project and the MITI's FGCS project of 1980s, I begin with a brief historical review of the amazing success of the Internet and its various applications and services that have been taking place in the past two decades.

I then discuss problems and difficulties that the ongoing evolutionary changes and enhancements of the present Internet will invariably face sometime soon. I give a summary of various rationales and motivations behind the recent initiative by the NICT's AKARI project and similar efforts launched in the U.S. and Europe within the past few years.

I will then discuss the opportunities and challenges that the NICT and Japan may face, using the AKARI Project as an example. I discuss these challenges in three parts: technical issues, management issues and political issues.

1. Introduction

Dear Honorable Senior Vice-Minister, President Miyahara, Director Enami, distinguished guests, and ladies and gentlemen. It's my great pleasure and honor to be invited to give this keynote address.

First of all, I would like to extend my heartfelt congratulations to those who have worked so hard towards the establishment of the Keihanna Research Laboratories. It is my understanding that research activities of this Keihanna Research Laboratories are primarily concerned with advanced applications such as natural language processing and cognitive information processing. These technologies will be increasingly important in the future, as we foresee ubiquitous use of information and communication technologies, which will profoundly transform the way we lead our life, conduct our business, and interact with the rest of the world.

I would like to focus my talk on a future generation network, which is referred to here in Japan as the “New-Generation” Network (NWGN). This particular term is used to be distinct from a similar term “Next-Generation” Network (NGN), which is currently pursued world wide as an extension to the existing Internet. The term NWGN should be also distinguished from another similar term “Next-Generation” Internet (NGI), which often refers to the currently deployed effort to migrate from IPv4 (Internet Protocol version 4) to IPv6 (Internet protocol version 6).

The NICT is going to play a central role in Japan’s efforts on NWGN by coordinating the AKARI project launched last November to cooperate, and in some sense to compete, with similar efforts that had started in the U.S., Europe and elsewhere. I have been recently appointed by President Miyahara to assist the strategic planning of the project, and I am looking forward to getting involved in this exciting effort.

It is rather difficult for me to visualize at this point exactly what such a new-generation network architecture will look like, and how new applications and services offered by such future network will affect our daily life, our society and the world. But I have some personal opinions and ideas on how this type of national project ought to be managed. You may find that my speech is rather more critical than is customary in Japan, especially on this type of occasion where we celebrate the start of the new Laboratories. But I hope that my comments on the Japanese efforts in R&D be taken as constructive criticism, not as negative and discouraging remarks. People in Japan may not recognize the problems I see. Even if you do, you may not be in a position to be able to speak candidly. Some of my American or European colleagues who are well versed in Japanese R&D environments may see the problems, but they probably dare not speak as straight from the shoulder as I can, because they would be afraid of offending you.

I am not an historian, so I am not interested in going back to the early history of computing and communications. It will be meaningful, however, to review the last twenty or thirty years, since many of us here have witnessed the changes and advances made in information technologies and emerging new applications during these periods. It is also true, however, that some segments of information technologies have not fulfilled our expectations. For instance, I don’t think that the field of artificial intelligence (AI) has made any significant advance in the past few decades.

Ever since I left Japan for the U.S. in 1965 to study at Princeton, I have primarily stayed in the U.S. except for the four years from April 1982 to June 1986 as the founding director of the IBM Japan Science Institute (JSI), later renamed as the IBM Tokyo Research Lab. It is not a total coincidence that the creation of JSI took place when the MITI (Ministry of International Trade and Industry)’s “Fifth Generation Computer System (FGCS)” project started. The FGCS was a very ambitious project with the goal of developing a *highly parallel processing* computer running on massive databases or *knowledge bases* using a *logic programming* language, *Prolog*. One of the main anticipated deliverables was to develop a prototype of a *high*

performance workstation. Applications of such powerful computers the FGCS team envisioned included an *expert system* built on an *inference machine* with advanced man-machine interfaces such as natural language processing, machine vision, and pattern recognition. Some of the research projects going at this Keihanna Research Laboratories may be traced back to application technologies planned in the FGCS project 25 years ago.

The FGCS project generated a great deal of excitement in Japan and some apprehension in the U.S. and Europe, which witnessed the Japanese dominance in the consumer electronics field since the 1970s and also the Japanese automobile's steady penetration into the U.S. that started in the 1980s. The MCC (Microelectronics and Computer Technology Corp.) was set up in the U.S. and ESPRIT (European Strategic Program of Research in Information Technology) in Europe. Despite its fanfare and big splash, the FGCS Project that lasted about 10 years was a big failure in my opinion. Today we seldom hear or see the name FGCS even in books on computer architecture. To be fair to my Japanese colleague, I should mention that Japan is not alone in colossal failures. Early HDTV effort in the U.S. (circa 1990) tried to leap frog groundbreaking Japanese work, but ultimately never became relevant. I'm sure there are more than enough European projects with similar fates.

It has been probably a taboo in Japan to talk about the failure of the FGCS, and it is not particularly constructive or pleasant to dwell on the negative side of the project that excited young computer scientists of Japan a quarter century ago. Today, Japan's economy is not as strong or threatening as 25 years ago, and we Japanese are not as confident as we were then. The AKARI project is perhaps much more modest and less ambitious than the FGCS was and is not getting such a monstrous attention here and abroad as the FGCS got. Furthermore, the players are different, and the project objectives are different in these two projects, yet it is my observation that Japan's culture, mentality and management style have not changed much in these past twenty five years. The failure of the FGCS project was partially attributable to the fact that the project team was dominated by the members of ETL (Electro-Technical Laboratory or "Densoken") of MITI, and opinions of researchers from participating companies were not sufficiently incorporated. Consequently, the prototype workstations or inference machines satisfied some AI researchers, but ended up with being "empty boxes" with no useful applications built upon them.

Thus, I cannot help pointing out some similarity between the FGCS project of the 1980s and the AKARI project of today. Both projects started at points in time when Japan has caught up with the U.S. in the current generation system, and Japan aspires to play a leadership role in the new generation system. The key strategy includes a *prototype* or a *testbed* of a new-generation system along with its architectural design. The major challenge in carrying out such a project is *how to consider applications from the early stage* of architectural design, *how to balance international competition and cooperation*, and *how to manage such well orchestrated project*. Thus, in my view it is important to understand where the FGCS project went wrong and to learn a lesson from that failure.

2. A Brief History of the Internet

As many of you may know, the origin of the Internet goes back to 1969, almost 40 years ago, when the Defense Advanced Research Project Agency (DARPA) of the U.S. Department of Defense commissioned a large wide-area network ARPANET (Advanced Research Project Agency NETwork), which served as a *testbed* for packet-switched technologies. The ARPANET, developed under the leadership of principal architect Larry Roberts linked many universities and national laboratories across the United States. The transmission control protocol (TCP), designed by Vinton Cerf and Robert Kahn in 1973 (published in IEEE Tran. on Communications in 1974 [1]), was introduced in 1977 for cross-network connections, and it began to replace network control protocol (NCP) which had been used in the original ARPANET. The TCP protocol was faster, easier to use, and less expensive to implement than NCP. In 1978 the TCP was separated into an *internet layer protocol* (IP) and a *transport layer protocol* (TCP), where IP (Internet protocol) is responsible for the routing of packets within the network that has limited intelligence. All that a router at a given network node does is to read the destination address (IP address) of an individual packet and forward it to one of its neighboring nodes. The functions that put together separately received packets into a correct order and deliver them to the end user (or application) rests with the *end system* of the network and TCP specifies these functions and services.

The protocol suite, TCP and IP, has since then been commonly known as TCP/IP. By the beginning of 1983, every site within or connected to the ARPANET switched to TCP/IP. From then on, all networks that use TCP/IP are collectively known as the Internet. The standardization of TCP/IP allowed the number of Internet sites and users to grow exponentially.

This sequence of events in the 1970s and early 1980s was probably well predicted by Vint Cerf and Bob Kahn, who are the recipients of this year's Japan Prize. The fact that it took only ten years from the conception of the TCP/IP until its deployment at all sites connected to ARPANET must have to do with the fact that, besides TCP/IP's superiority, the co-inventor Dr. Robert Kahn himself served as Director of DARPA's Information Processing Techniques Office (IPTO) and was responsible for the ARPANET project. Drs. Cerf and Kahn are both visionary individuals, but I do not think even they could foresee the enormous impact that the ARPANET and their TCP/IP protocol have brought to our life, businesses and societies, as we witness today.

The original objective of the ARPANET was merely to allow efficient sharing of programs, files and data among researchers at different universities and national laboratories in the U.S. The first unexpected surprise was that the network traffic was dominated by email rather than by file transfers. The ARPANET became the NSFNet in 1985, using 56 kbps backbones, which was subsequently up-scaled to 1.5 Mbps in 1986.

The public interest in the Internet accelerated once the commercial email service of the Internet was introduced in 1989, and three Internet service providers (ISPs) got started. The World Wide Web was invented by Tim-Berners Lee at CERN (European Particle Physics Laboratory in Geneva, Switzerland) in 1989, and its use expanded in the 1990s with the advent of web browsers such as Mosaic, Netscape and Internet Explorer.

Amazon.com, Inc. founded by the then 30 year old Jeff Bezos (a Princeton EECS graduate) in 1994 was among the first dot com companies. Yahoo and Alta Vista were founded both in 1995. In 1999 Google was founded by two Stanford Ph.D. students Larry Page and Sergey Brin. Their search algorithm, PageRank, is different from the previous search algorithms and is based on probabilities associated with a Markov chain, where the states are pages. In 1997 a Canadian based company Research in Motion (RIM) introduced a wireless handheld device BlackBerry. YouTube, LLC (UTube), which provides a website where users can upload, view and share video clips, was founded in 2005 and was bought by Google in 2006. Although Apple introduced in 2007 "iPhone", an Internet-enabled multimedia mobile phone, mobile phones are an area that the Japanese are far ahead of the U.S.

3. The Internet Deployment in Japan.

In 1991-92 I took a sabbatical leave from Princeton and was appointed by the late Professor Takanori Okoshi to the NEC C&C chair visiting professor at the RCAST (Research Center for Advance Science and Technology), the University of Tokyo. My recollection was that there was little awareness and interest in the Internet in Japan. As a matter of fact, at that time the Broadband ISDN (B-ISDN), which was based on ATM (asynchronous transfer mode) fast-packet switching (also known as cell relay) technology, was touted by telecommunications carriers (AT&T, NTT and European PTTs and others) as a unified network to provide broadband (i.e., high-speed) multimedia services for the 21st century.

I am yet to find out what actually happened in the 1990s here in Japan, but if I remember correctly, the bandwidth cost or tariff in Japan was much higher than in the U.S. and South Korea in the early 1990s. But this situation seems to have been reversed in the past several years. Although NTT was rather slow in recognizing the significance and threat of the Internet development, its longstanding policies since the 1970s to develop all digital high-speed network infrastructures should be commended. In 1979, NTT announced its INS (Information Network Service, all-digitalization of networks) deployment plan. In the mid 1980s, the Japanese government decided to pursue policies of liberalization that opened Japan's telecommunication market to competition and moved toward the privatization of NTT. In 1990 NTT introduced its VI&P (Visual, Intelligent & Personal) deployment plan and in 1993 introduced its FTTH (Fiber ToThe Home) deployment plan. This series of deployment plans towards high-speed multimedia services provided necessary infrastructures to support the successful deployment of high-speed Internet services that ensued.

It is noteworthy that in 1999 NTT DoCoMo, which was spun off from NTT in 1991, introduced “i-mode” service, a wireless Internet service using cellular phones. The huge success of i-mode service has a lot to do with the proliferation of the Internet among the population in Japan. In the world scene, however, i-mode is virtually unknown [2].

It is my understanding that the following developments have contributed to the explosive growth in Internet users, advanced Internet services, and the lowered tariff in high-speed services in Japan.

First, the Ministry of Internal Affairs and Communications (MIC), with a strong push from the industry, launched the “e-Japan Program” (2001-2004) as a national strategy to catch up with the U.S. and other nations in the deployment of information and communication technologies. During the short period of three and a half years from March 2001 through August 2004, the number of high-speed Internet subscribers increased twenty-fold from 0.85 million to 17 million, according to a report released by the MIC [3]. More than 46 million households became DSL (digital subscriber line) subscribers and 36 million households were connected to optical networks through FTTH (fiber to the home). Both numbers were significantly higher than the original goals set in 2001 (i.e., 30 million households for DSL and 10 million households for FTTH) [4]. However, I have not been able to confirm the validity of these numbers. I find conflicting numbers in the Telecommunication Industry News [4] of July 14, 2007. It reported that at the end of March 2007, Japan had a total of 26.4 million broadband subscribers, 33% of which use FTTH connections. A majority of customers (63%) continued to use DSL internet services, the report says. Another conflicting report is found in Webmaster World.com [5] which stated on April 5, 2004 that at the end of February 2004, there were 1.04 million FTTH subscribers, 10.9 million ADSL subscribers and 2.55 million Internet cable subscribers. As of this writing, I cannot resolve this huge disparity of the statistics between [3] and [5]. They are off by a factor of 4.5 for DSL and by factor of 35 for FTTH subscribers as of 2004. Either one or both of the reports must be grossly incorrect!!

The MIC launched in 2004 the “u-Japan program” (2004-present) with the goal of making Japan the most advanced nation in terms of “ubiquitous” use of the information technology. Its specific strategies includes the migration to the IPv6 (Internet Protocol version 6) from the current IPv4. IPv4 uses 32 bits to represent an address, thus the network can have at most $2^{32} \cong 4$ billion different addresses (cf. There are roughly 6.67 billion people in the world [6]). IPv6 uses 128 bits to represent an address, and $2^{128} > 3.4 \times 10^{38}$, which is, practically speaking, virtually inexhaustible. This means we can assign unique address or ID to every item around us, not only electronic devices but virtually all things around us, if we want. This will have enormous implications in our business and daily activities in our future society, if the computation, communication, and storage capacity of future information networks can handle such vast information. We should be able to keep track of everything in

this world if we have a way to tag its ID on it, and collect, transport and process the information concerning their locations and status by using some sensing mechanisms.

The most recent program initiated by the MIC is the formation of “Ubiquitous districts,” where advanced new ICT (information and communication technology) services will be tested. I am much interested in finding out how this experiment is progressing.

4. Evolutional changes of the Internet

The Internet that started in the 1970s have gone through a series of innovative changes and enhancements. As I stated in Section 2, TCP/IP was originally designed when the main purpose of the Internet was to share files and data among researchers at universities and national laboratories in the U.S. The notion of QoS (quality of service) was absent and the issue of security not considered. Real-time applications such as voice services and time-sensitive applications such as video streaming were not even envisioned. At that time virtually all communication links were fixed wired lines, and terminals were fixed terminals: although ALOHANET at the University of Hawaii and a ground-radio packet switching network [7] were part of the ARPANET project, mobility of the terminals was not a major consideration.

The basic philosophy behind the original Internet architecture is that all information (originally just programs, files and data, but now all sorts of information including, voice, image, video as well) is chopped into smaller pieces and encapsulated to form entities called “packets” (typically several hundreds to a few thousand bits long). Each packet carries the address of its destination node (IP address) in its “header”, and the router at a given node reads the address and forwards the packet to one of its neighboring nodes. This is essentially what the internet (IP) layer does and all data packets are treated *equally*. *Fairness* is the basic philosophy in handling data packets when congestion occurs in some part of the network. The service provided by the IP is rudimentary and is called a *datagram* service: arriving packets at the destination end system are not necessarily in the right sequence, and some of the data packets may even be missing. It is the responsibility of the end system to collect all the packets and deliver them to the end user or the application in the right order. This high-level function is the responsibility of the transport layer, and the TCP handles this end-to-end service.

In the past fifteen years or so, especially after the Internet has been used for commercial applications, a series of extensions and modifications have been made to the original TCP/IP protocol suite. They include a feature to perform network address translation (NAT) that enables a local area network (LAN) to use one set of IP addresses for internal traffic and a second set of addresses for external traffic; to differentiate the quality of service (QoS) depending on the importance of applications/users (DiffServ); to provide security for data packets routed within the network (IPSEC), and to control access to an edge node (Firewalls); to provide voice services over the Internet (Voice over IP or VoIP); to support streams of video

information (video-streaming); and to allow mobile devices/users to move from one network to another while maintaining a permanent IP address (mobile IP).

All these protocol designs have been done by the standards group called IETF (Internet Engineering Task Force), which started in 1986 as a U.S. government funded activity, and is now an independent, international activity associated with the Internet Society (ISOC) formed in 1992. Incidentally, the ISOC's Japan Chapter was established in 1994. In parallel with this group, the Internet Association of Japan (IAJ) was formed in 1992 and was formally recognized by the MIC and METI in 2001.

It is amazing that the 35 year old protocol TCP/IP has remained the core of the Internet protocol, and the currently ongoing international effort called NGN (Next-Generation Network) will continue to support the architectural evolution around the Internet Protocol (IP) for both core and access networks over the next five to ten years.

In parallel with these engineering efforts to enhance the IP-based network, there has been an increasing level of concern and discussion in the networking research community since around 2000 (see e.g. [8]) as to how long these incremental changes or "patchwork" can continue to be a viable approach to meet our future needs. Despite the various enhancements made to the Internet mentioned above, it seems increasingly difficult to deal with the complexity of the Internet that has grown over the years, which will ultimately reach the point where the maintenance and error recovery will be exceedingly difficult. One of the few quantitative arguments that indicate the limit of the Internet is made from the view point of electric power consumption of the Internet routers ([9] page 7). Although it seems rather difficult to prove the validity of these quantitative forecasts of power consumption, the complexity argument alone seems sufficient to justify a completely fresh look at a network architecture required for our future. There seems to be a strong camp in the networking community which advocates that this "clean-slate" approach ought to be investigated without being constrained by the IP based architecture. That is to say, so-called "backward compatibility" should not be given high priority when we explore a new architecture. They argue that the issue of compatibility with, or migration from, the IP network ought to be looked at only after the new network is designed.

5. What should the NWGN be like?

The excellent document [10] prepared by the AKARI project team in April 2007 is 120 pages long and is quite comprehensive. What I present below is not a summary of their document, but rather it is a list of several requirements of such a future network as I see it. The order of these requirements is not necessarily in accordance with their importance.

- a. Must support a large number of ubiquitous and different applications. In other words, *scaling* and *heterogeneity* of the network will be major attributes.

- b. Must support highly dynamic mobility in rapidly changing networks. In other words, mobile users and time-varying network resources must be handled.
- c. Must be flexible enough to deliver services envisioned (but not necessarily well defined) for future societies and business practices.
- d. Must allow energy efficient implementation. All-optical processing and “virtualization” of network resources seem to be promising technologies.
- e. Must be robust and secure against network failures, malfunctions, and malicious and sophisticated attacks.
- f. Must be flexible and open to allow support for unforeseen implementation technologies and unforeseen applications.

In 2006 the U.S. National Science Foundation (NSF) launched a new funding program by inviting the research community to consider what the requirements should be for a global network of 15 years from now. The intellectual scope of the FIND program is listed as follows [11]:

- How can we design a network that is fundamentally more secure and available than today’s Internet? How would we conceive the security problem if we could start from scratch?
- How might such functions as information dissemination, location management or identity management best fit into a new network architecture?
- What will be the long-term impact of new technologies such as advanced wireless and optics?
- How will economics and technology interact to shape the overall design of a future network?
- How do we design a network that preserves a free and open society?

In parallel with the FIND program, the NSF has been sponsoring a prototype/testbed project called the GENI (Global Environment for Network Innovations) Initiative, which will allow different architectural ideas to be tested [12].

In Europe, the EuroNF project [13] is funded by the European Commission in FP7 (the Seventh Framework Programme) as a follow-on of the EuroNGI/EuroFGI. It is a Network of Excellence on the Network of the Future, formed by 35 institutions (the academia and industry) from 16 countries, and its main target is to integrate the research effort of the partners to be a source of innovation and a think tank on possible scientific, technological and socio-economic trajectories towards the network of the future. Its scientific coordinator is Dr. Daniel Kofman, a professor of ENST, Paris and the CTO of RAD Data Communications, a Haifa-based company.

6. Opportunities for the NICT and Japan

The NICT is in a unique position to coordinate Japan’s nationwide effort on the architectural design and testbed construction for beyond 2020. To the best of my knowledge, neither the U.S. nor the European group has a funding agency that is equipped with an in-house research lab. The AKARI project has an additional advantage in that the NICT has in-house expertise in advanced and emerging

technologies, especially wireless/radio communication technologies at the Yokosuka Center and “universal” communication technologies here at the Keihanna Laboratories.

The NSF, however, announced in May 2007 that BBN Technologies (formerly, Bolt Beranek and Newman) was selected to serve as the GENI Project Office (GPO) with Mr. Chip Elliot, Chief Engineer of BBN, as the project director of GENI. BBN has a reputable track record by having served as the implementer of the ARPANET (Kahn and Cerf worked for BBN briefly during the development of the ARPANET).

The AKARI Project of the NICT got started last year, not much later than the announcement of the FIND (Future Internet Design) program by NSF and EuroFGI, a similar program in Europe, now called EuroNF. I anticipate that Japan’s advanced communication infrastructures and high-speed Internet applications that have been taking place in recent years will provide important leverage to the AKARI Project over its competitors, because the architectural designers will benefit from experiences and performance data to be provided by these advanced applications.

Although Japan had a slow start in joining the Internet evolution as stated in Section 3, Japan has now the most advanced cellular system deployments and applications (such as DoCoMo’s “i-mode” services) in the world. The Japanese government and industry have in recent years been pushing hard the slogan of “an ubiquitous IT society” to be built on the Internet and wireless networks. As a result, Japan has developed many leading edge applications such as text messaging over a cellular phone, smart highways and smart cars, intelligent homes, IT for “smart” assisted living, and advanced computer games. Thus, as the cellular and the Internet merge into a new generation network (NWGN) architecture, there will be tremendous opportunities for Japan’s further creativity in both wireless core technology and ubiquitous applications. Prof. Raychaudhuri, Director of WINLAB of Rutgers University, who is well familiar with Japan’s advanced technology and applications in wireless networks, thinks that the NWGN project in Japan could lead the way in developing unique pervasive/ubiquitous applications. He suggests that as much as 50% of the project resources and efforts should be devoted towards developing applications and end-user devices. As I pointed out earlier in reference to the FGCS project, it is important to select an architecture with ease of application developments in mind. The GENI/FIND programs in the U.S. have been broadly criticized for the lack of applications perspective, and this might be an area where Japan can excel, he says. In terms of core network technology itself, he thinks that the movement towards open networks and cellular-Internet convergence will serve to Japan’s advantage, provided the R&D groups in Japan can successfully adopt more competition-based open standards and an open source mindset, which has so far been lacking in the proprietary patent-oriented culture of Japan. This strategy should also help to involve more young Japanese people outside of major R&D organizations.

7. Challenges for the NICT and Japan

The AKARI project team made an excellent start by generating a comprehensive planning document a year ago, which covers a broad range of component technologies of the future, anticipated societal needs and corresponding design requirements as well as forecasted network traffic. It also discusses candidate network components, such as optical processing, radio access networks, algorithms for transport, addressing, routing, security and QoS. I plan to spend the next few months to digest this valuable document thoroughly and discuss with the authors of individual sections so that I can grasp the whole picture of the directions and status of the project so that I can serve effectively for the project.

With the risk of making some irrelevant or even wrong remarks at this point, I would like to express my opinions and some concerns so that I can provide some useful advice for those who are involved in this exciting and challenging project. I divide my discussions into three parts: (i) technical issues, (ii) project management issues, and (iii) political issues.

Technical issues:

Perhaps the most difficult aspect in conducting this type of exploratory research project is how to strike a balance between “focusing” and “being open.” The AKARI project team must have discussed up to now a number of plausible ideas that are worth exploring, and I presume that the team is now in the stage where its efforts are to converge towards a yet to be determined AKARI architecture.

The Japanese R&D team has traditionally been effective and productive in reaching a common target in an orchestrated effort, as long as the goal and direction are clearly identified. This was certainly the case for all national computer projects prior to the aforementioned FGCS project. The goal then was to catch up to the most advanced computers developed by IBM. I suppose that the situation was similar when the MIC launched the “e-Japan” and “u-Japan” programs at the beginning of this decade, i.e., to catch up with the rest of the world in participating in the Internet evolution.

This strength of the Japanese R&D style, however, can be its weakness at the same time, if the project effort gets narrowly focused too soon and excludes alternative ideas and suggestions that may come along during the execution of the project. That was certainly the case with the FGCS project. In my observation, the FGCS team reached a consensus in the early phase of the project, firmly convinced that a highly parallel processing architecture with massive databases and software written in logic programming was clearly the way to go. Furthermore, they were oversold by some AI (artificial intelligence) researchers in the U.S. and elsewhere who needed a boost for their technical community, which had been languishing for some time.

I imagine that there will invariably be some pressure from the top management and funding sources that a working prototype should be up and running by the end of the project period. Thus, Japanese research initiatives tend to quickly adopt a single architectural approach (as in the FGCS project), and after that there is a great deal of

focus on high-quality implementation of the adopted system. Such an approach can be very dangerous for the NWGN project, where the temptation will be to quickly agree upon a design that meets the requirements listed in the white paper, so cautions Prof. Raychaudhuri. The only way to avoid this type of mistake is to have multiple competing groups with relatively independent funding, and not to enforce a premature technical consensus among these groups. He suggests that the risk of a mega project failure can be substantially reduced by allowing for competition, having open standards and mostly open source software results, open platforms (like GENI of the NSF). Such strategies will enable smaller research groups to participate, avoiding early stage system integration, but instead focusing on demonstrating key architectural principles and delivering enabling system components.

I believe that it is of the utmost importance for the AKARI team to stay alert and to be attentive enough to observe how emerging enabling technologies might impact candidate architectures and to find out what other researchers are doing, while focusing on the direction the team has “tentatively” set. The NICT and its alliance members should keep the AKARI project as open to other research groups, both in Japan and abroad, as possible. Overconfidence is more dangerous than the lack of confidence. Be aware that there is almost always someone out there who is a bit smarter than we are.

Project management issues:

A good manager or leader must have four attributes. First, he/she must be technically competent enough to understand the work of his/her people and provide some directions for them. Second, the manager should be able to communicate effectively with his/her peers and upper management, as well as with his/her own people. Third, the manager must be sensible and sensitive enough to understand the needs of his/her team members. Some may require frequent positive strokes, while for others candid but constructive criticisms may be warranted. Fourth, which is perhaps the most important, is that a leader should respond to inquiries, proposals and questions promptly. If you cannot make a decision or suggestion on the spot, at least acknowledge that you have received the inquiry. A slow and indecisive manager demoralizes his or her people, and slows down the work process of the entire organization. If you act promptly, your people will also respond quickly and diligently. It is human nature. After all, “time” is the most precious resource for any productive individual or organization, and managers ought to act accordingly.

It is always difficult to find qualified managers, especially in a situation where available slots for recruitment from outside are limited. The NICT and other organizations like AIST (National Institute of Advanced Science and Technology) of the METI (Ministry of Economy, Trade and Industry) nowadays appoint university professors as part-time group leaders or advisors, and recruit young researchers as post doctoral fellows. I think these are excellent solutions, given the headcount constraint at the NICT and other national laboratories.

I should call to your attention that the current job market in the area of networking and communications is not so good in the U.S. Neither Bell Labs (now under Alcatel-Lucent) nor AT&T Labs is actively hiring, and teaching positions at good universities are extremely difficult to come by. So the NICT and other laboratories in Japan should be able to attract really bright people, if we are determined to do so. We need to place advertisements in appropriate places (e.g., IEEE Spectrum magazine) and send out announcements by email to the right professional communities (e.g., IEEE's computer communications committee, ACM's SIGCOM) and by directly contacting key professors.

Associate or full professors in the U.S. universities are typically granted a sabbatical leave of six months (with full salary) to a year (with half salary) every seventh year. So if we contact well in advance professors whom you want to have as visitors, there is a reasonable chance to get some of these people for several months to a year. Of course, professors who carry many research programs and lead a large research group may not be easily movable.

We, the Japanese, have traditionally looked up to the U.S. and Europe for first-rate researchers and novel ideas, but the world map of brains and innovations is rapidly changing. Now we find an increasing number of creative people in such nations as Australia, Canada, China, Hong Kong, India, Korea and Singapore. A fair number of faculty members at top-rated universities in these countries earned their doctorate degrees at reputable universities in the U.S. and U.K. These people have sought academic positions there (quite often their home countries) primarily because tenure-track positions in leading U.S. and European universities are extremely scarce in recent years. Singapore is especially aggressive in attracting brilliant talent to their universities and national laboratories (e.g., A*Star). They are also sending their talented young students to Ph.D. programs at universities within Singapore as well as in the U.S. and U.K. on a large scale (a few hundred Ph.D. candidates per year) with full support (i.e., tuition and a stipend or salary enough to live in the country where the student studies) [14].

Japan has been doing a terribly poor job in this regard, by not encouraging our brightest students to pursue Ph.D. programs here or abroad. Japan fails to attract top quality students and faculty members from outside of Japan. In all major research universities in the U.S., every admitted Ph.D. candidate in science and engineering is exempted from the tuition (\$35,400/year at Princeton) and receives a stipend (\$30,000/year at Princeton). This applies equally to both U.S. and foreign students.

Coming back to the issue of management, I am aware that the NICT appoints an external review board who examines both research programs and management policies of the labs. I participated in such review meetings twice in the past. It is an excellent practice, but I suggest that the review board should be a team of international experts in the case of the AKARI project. This is a most effective scheme whereby the NICT and AKARI project can have international exposure and get useful information about similar efforts and people in other countries.

The review panel members should be technical managers or professors who are actually involved in new generation network projects or similar research programs in their own countries. The review should be “real working sessions,” in which the panel members freely interact with group leaders/managers and researchers. Only through such an in-depth discussion and Q&A sessions, will the participating researchers and managers receive useful feedback, criticisms and recommendations. Top management and administrative staff should meet with the review panel at the end of their visit, receiving their assessments and recommendations.

A high-level “show and tell” review may serve well for public relations of the laboratory or the center, so you have to do it once in a while. But the researchers and their leaders who prepare for presentations or lab tours seldom receive useful feedback or suggestions, because the members of a high-level review board are typically directors of some laboratories or VPs of companies or senior professors (like me), and they are seldom really engaged in research activities themselves. So all you can expect to get from these folks are comments like “Very interesting and nice work!” or “Keep up the good work!” Many Japanese universities and national laboratories nowadays have so-called external-evaluation committees, but to my knowledge none or very few, if any, has implemented the kind of review meeting I have in mind.

Any review meeting from which we cannot expect useful feedback will be a waste of time, and hence should be minimized. As I said earlier, time is the most precious resource for any productive organization.

Political issues:

It is my understanding that the NICT set up the “NWGN R&D Strategic HQ” last October with President Miyahara at the helm, so this group might have already given ample thought to the questions and suggestions I raise below.

It is not clear to me at this point to what extent the research strategies and results of the AKARI project should be treated as proprietary information. Considering that the huge success of the Internet owed, to a great extent, to the fact that its development and operation were closely tied to those of Unix, which provided basically free and open resources. Everyone enjoyed free access to the Internet, the Unix operating system and related software as long as he/she was connected to a node that is recognized by the Internet community. Given this culture and tradition of the networking community and the intrinsic nature of networking, in which compatibility and open architecture are essential, we ought to be quite open in exchanging ideas with our counterparts in the U.S. and elsewhere. It will be worthwhile, however, for your IP (intellectual property) experts to look into this matter, by reviewing for instance the NSF’s policies on IP matters and their rules concerning participations by researchers in foreign countries.

When it comes to business and commercial applications, however, it is quite a different matter. For instance, I presume that Google must have a great stake in their patents on their search engine algorithms and RIM (Research in Motion) must hold some key patents on their BlackBerry products and services. As is well known, in the case of Qualcomm, IP licensing fees are now their whole line of business. In this sense, it is essential for the NICT and its industrial partners to take a leadership role in the research, development and standardization efforts of a new generation network (NWGN) architecture.

Technology transfer is often the most challenging step in turning your idea to reality. In the case of network architecture research, its ultimate technology transfer is to convince the international community to accept your architecture as an international standard, and help the industry to produce cost-effective products and related services. The technology transfer of your architecture (whether it be a computer architecture or a network architecture) is likely to be met with resistance by other contenders, because of the so-called NIH (not invented here) syndrome. No individuals, companies or groups are unwilling to accept your architecture unless your architecture can be proven to be overwhelmingly superior to theirs. In other words, you must back up your claim with some objective data such as performance and/or cost advantage.

The ability to argue effectively and persuade other parties will be often important in technology transfer negotiations. I presume that the same is true for negotiations at standard committee meetings. According to Dr. Gerhard Fasol of Eurotechnology Japan, K.K., several Japanese inventions failed to make global impacts, because the Japanese engineers could not impress other members of the international standardization committee. Dr. Akira Sekino, formerly with NEC Corp., suggests that it is strategically important to establish close communications and intimate relations not only with network and computer makers but also with semiconductor manufacturers, because network functions that require high-speed processing must be implemented in hardware. The timely development of such chip-sets is crucial to obtain strong support in the standardization committee meetings.

The younger generation Japanese are more fluent in English and have a better command of spoken English than my generation. Well prepared PowerPoint slides may overcome some deficiencies in written texts or oral presentations. But if your written sentences are full of grammatical errors, it is quite natural that others may undervalue your technical competence and hence may not take your work or proposal seriously. A group like the AKARI project should hire a capable English editor who can help editing reports and manuscripts written in English. He or she should be also asked to help oral presentations of AKARI members through rehearsals. The IBM Tokyo Research Laboratory has had such a person in place for sometime.

8. Concluding Remarks

The greatest challenge for the “clean-slate” architecture like the AKARI project will

be perhaps resistance from those who continue to push for evolutionary changes of the existing IP-based architecture. Unless we are fortunate enough to come up with a viable migration path from the IP-network to the new generation network (NWGN) or can design an “emulator” that supports the IP-network applications, running on the NWGN, both the IP-network and the NWGN will have to coexist in some fashion for some period in the future. New applications and services may run on the NWGN, whereas the existing applications and services must continue to run on the IP-network. Then some clever way to “virtualize” network resources may be a workable solution.

I presented some technical, managerial and political challenges that the Japanese R&D groups may face in system projects that require “collective creativity.” Many of these challenges can be avoided or mitigated by recruiting capable foreign researchers and managers. The Japanese government, industries and academia are still very closed and inward looking. The number of foreign professors in major universities and the number of foreign researchers and engineers who are appointed managers and executives in Japanese organization is extremely small. Appointments of foreign nationals to key positions will be one of the effective means to strengthen the human capital in Japan.

Acknowledgments:

In preparing this talk, I asked for help from a number of friends. They are (in the alphabetical order) Dr. Toshiharu Aoki (NTT Data), Dr. Philip Chou (Microsoft Research), Dr. Gerhard Fasol (Eurotechnology Japan K.K.), Prof. Brian L. Mark (George Mason University), Prof. Masayuki Murata (Osaka University and AKARI Project), Prof. Yoshikuni Onozato (Gumma University), Prof. D. Raychaudhuri (WINLAB, Rutgers University), Dr. Akira Sekino (formerly NEC), and Prof. Phuoc Tran-Gia (University of Wurzburg, Germany). I am especially thankful to Prof. Raychaudhuri and Dr. Sekino for their valuable comments concerning my earlier draft of this text. Dr. Chou and Prof. Mark have read the entire manuscript very thoroughly and helped me with numerous suggestions to improve my English.

I have just received valuable comments and additional information from Prof. Jennifer Rexford of Princeton University regarding the GENI/FIND projects. I will incorporate her input in my final version, although I cannot do it for this semi-final version to be submitted to the conference organizer.

The criticisms about the FGCS and other Japanese efforts or products are entirely based on my observation and thought, and the above-mentioned friends have had no part in it.

I would very much appreciate receiving your comments or opinions for any part of this document. I am open to any suggestion or criticism to improve my ideas or opinions expressed in this document. I can be reached at hisashi@princeton.edu or at 347 West 57th Street, #33B, New York, NY 10019, USA. Phone/Fax: (212) 977-1897

References:

- [1] V. G. Cerf and R. E. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Trans. Commun.* Vol. 22, 1974, pp. 637-648.
- [2] Chritsn Caryl, "Why Apple Isn't Japanese: Once a technological leader, Japan is now struggling to find its place in the digital age," *Newsweek*, Dec 1, 2007
<http://www.businessweek.com/id/73236/>
- [3] Ministry of Internal Affairs and Communications, "Contributing to Future Deployment of e-Japan Strategies,"
http://www.soumu.go.jp/menu_02/ict/u-japan_en/new_outline01.html
- [4] Telecommunications Industry News, "FTTH Begins to Eclipse DSL in Japanese Broadband Market," July 14, 2007. <http://www.teclick.ca/2007/07/ftth-gegins-eclipse-dsl-in-japanese-broadband-market/>
- [5] WebmasterWorld.com, "Asia and Pacific Region," April 5, 2004.
<http://www.webmasterworld.com/forum32/605.htm>
- [6] U.S. Census Bureau, "U.S. and World Population Clocks – POPClocks,"
<http://www.census.gov/main/www/popclock.html>
- [7] L. Kleinrock, *Queueing Systems, Volume II: Computer Applications*, John Wiley & Sons, 1976.
- [8] R. Braden, D. Clark, S. Shenker and J. Wroclawski, "Developing a Next-Generation Internet Architecture", a White Paper, July 15, 2000.
<http://www.isi.edu/newarch/DOCUMENTS/WhitePaper.pdf>
- [9] 村田正幸, "新世代ネットワークに向けた研究開発動向," 2008年4月
- [10] **AKARI** プロジェクト, "新世代ネットワーク・アーキテクチャ **AKARI** 概念設計書," 2007年4月.
- [11] National Science Foundation, "FIND: NSF NeTS FIND Initiative,"
<http://www.nets-find.net/>
- [12] D. Raychaudhuri, "Future Internet Research in the US: NICT Meeting, Washington D.C., March 2008. See also www.winlab.rutgers.edu/WMPG
- [13] EuroNF Public Web Portal. http://euronf.enst.fr/en_accueil.html
- [14] Philip Yeo, "Trying to Build a Knowledge Based Economy," presented at a Special Colloquium at Princeton University, May 8, 2008.