4657 words / 32.1 minutes

APAN50 Virtual Meeting and Conference – Hong Kong
Keynote Address
Michael A. McRobbie
President, Indiana University
Tuesday, August 4, 2020
1:00 p.m. HONG KONG TIME
1:00 a.m. EDT

1. INTRODUCTION

Thank you very much, Professor Goto, for that kind introduction. Goto-san is an old friend whom I have known for many years, and he was a central part of the very first meetings and discussions that led to the establishment of APAN. He has had a long and distinguished career at Waseda University, and he continues to serve as an advisor to the APAN Board of Directors. A few years ago, Goto-san was—deservedly—inducted into the Internet Hall of Fame.

I am very pleased to take part—virtually—in this major historical milestone for APAN, the 50th APAN meeting—and the first to be held as a conference, attended as I understand it by around 1,000 people from over 50 countries. I know that many old friends and colleagues are with us, and I am delighted that two of the other key founders of APAN, Professor Kilnam Chon and Mr. Kazunori Konishi, are with us. I want to thank Dr. Louisa Lam, director of the Joint Universities Computing Center in Hong Kong and chair of the local organizing committee, and my former colleague, Markus Buckhorn, APAN's general manager, for their kind invitation to speak. Let me take this opportunity to also congratulate, on behalf of Indiana University, the Joint Universities Computing Center on its 50th anniversary. Such anniversaries are very important and symbolic in the history of an institution as we are particularly aware at IU as we conclude a year-long celebration of our 200th anniversary. So, I congratulate JUCC for its long history of inter-institutional collaboration and its role in the development of information technology in the Hong Kong higher education community.

I am very proud to have been one of the original founders of APAN and to have been a member of the global IT community for many years, both when I was in Australia and then in the United States as CIO at IU. During this time, I worked on many issues in IT and networking with people from all over the world. Like many of you, I had a certain perspective on these issues and their importance, but this was necessarily bounded by what were, relatively speaking, responsibilities focused in just one area of the much larger enterprise of the entire university.

But in 2007, I became president of one of America’s largest research universities with a budget now approaching $4 billion, nearly 100,000 students, and 40,000 faculty and staff, as well as a $6 billion hospital system. In this position my perspective rapidly evolved and broadened. Some of the issues that once seemed of enormous importance, that filled me and, I know, many other people who have joined us for this conference, with passionate intensity, seem of far less import when one is struggling with huge budget issues caused by government disinvestment in higher education, with damaging government policy changes, with an increasing mood of public skepticism about the value of science and education, with catastrophic potential risk ranging from athletics to infrastructure, and with existential threats to our very existence as institutions, such as the global pandemic that we are now all experiencing.

These are all, of course, more immediate challenges. But from a much longer-term perspective, a perspective not even of decades but of centuries, those of us who lead large research universities must also ensure that all that we do, including in IT and networking, reflects the three fundamental missions of a university: the creation of knowledge (that is, research and innovation), the dissemination of knowledge (that is, education and learning), and the preservation of knowledge (that is, information repositories).

So, in this address, I hope to address the theme of this conference – “the future of networked research and education” – by addressing some of the issues that I think are of long-term institutional importance and which transcend the daily battles we have to fight in higher education. But first, given the historic nature of this APAN meeting, I want to say just a little about the formation of APAN in 1996 from someone who was quite literally “in the room where it happened.”

2. APAN PRE-HISTORY AND HISTORY

My own interest in high-performance networking was very much influenced by the work done on the US Gigabit Testbed Initiative (GTI), which commenced in about 1990 and ran for about five years. There were five testbeds, but the one I believe was the most suggestive and visionary was CASA—a gigabit testbed that connected Caltech, JPL, the San Diego Supercomputer Center and Los Alamos National Laboratory. The primary focus of CASA, under the extraordinary leadership of Dr. Paul Messina, was on distributed heterogeneous supercomputing applications involving very large computational problems. As such, it was a very influential forerunner to grid computing and much more. The GTI, in part, led to the establishment of the very high-speed Backbone Network Service (vBNS) and, indirectly, to Internet2.

The GTI, and CASA in particular, influenced the establishment of a new research center in Australia of whose board I was a member—the Research Data Network Cooperative Research Center (RDN-CRC), which was established in 1993 to carry out research on Telstra Australia’s—Australia’s major carrier—Experimental Broadband Network (EBN). Two of the research projects were in heterogeneous supercomputing.

Around the same time, discussions were being carried out between the Australian and Japanese governments—they had even risen on two occasions to the prime ministerial level—on cooperation in high-performance computing that could broadly build on the success of the research partnership between the Australian National University and

Fujitsu, in which I was then involved, and which was so ably led for many years by Professor Robin Stanton, who I believe is attending this conference. These discussions led ultimately, in 1995, to an initiative to establish a direct network connection essentially to connect high-performance network testbeds in Australia and Japan.

Through these activities, I was invited to give a paper at the *Asia-Pacific Economic Cooperation Symposium for Realizing the Information Society* conference, which was held in Tsukuba, Japan on March 27-28, 1996. The paper was called "Towards an Asia-Pacific Gigabit Testbed." In it, I advocated generalizing the model of intercontinental connectivity that we were trying to bring about between Australia and Japan, to the broader Asia Pacific regions based on the success of GTI in the United States.

Professor Kilnam Chon, then in the Department of Computer Science at the Korean Advanced Institute of Science and Technology, also gave a paper at the same meeting called "Asian Highway – Why Not Gigabit Network?" Though we had never met, nor did we know each other at that point, our papers were remarkably similar in what we advocated—namely establishing a high-performance Asia Pacific network.

During the meeting, a small group of us—including myself, Professor Chon, Professor Goto, Mr. Konishi, then at KDDI, Dr. Tin Tan Wee from Singapore, and one or two others—met in the lobby of the Dai-Ichi Hotel in Tsukuba on the evening of March 27, and agreed to try to establish an initiative to do this.[1] This initiative proposed the establishment of what we then called the AP Testbed, and I drafted the summary of our meeting (which I still have).

Steve Goldstein, then at the National Science Foundation, also attended the meeting just to observe. He did, though, inform us that the NSF would be releasing a solicitation for the High-Performance International Internet Services (HPIIS) program, calling for

---

[1] I am told that the couch where we met for that discussion is still there, and is sometimes pointed out with reverence by members of the APAN community as the place where the organization had its inception.

proposals to establish connections between the then vBNS and comparable networks in other regions of the world. He suggested that this program might be a way of providing a high-speed connection from the Asia Pacific to the United States.

On June 18-20, 1996, the *APEC/APII Testbed Forum* meeting was held in Seoul, Korea. By this time, Professor Chon and I had agreed to combine and revise our papers into a joint paper called “Towards an Asia Pacific Advanced Network,” which we presented at this meeting and which was where the name APAN first appeared.

During the *Forum,* we had the first major meeting of representatives of most of the countries, institutions, and networks that would eventually become the founding partners in APAN. At this meeting, we had a discussion in more detail about what would be APAN’s mission. Would it be, in my words, more of a coordinating body for regional advanced networking activities, or would it be an organization more like, for example, Internet2, that owned and operated its own infrastructure funded by its members that interconnected all of them, as well as providing connectivity to other nation’s research networks?

But the latter highly ambitious option was always going to be difficult to accomplish, as it was an idea well beyond what most funding agencies could even comprehend at that point. Recognizing this reality, the meeting agreed that APAN’s mission would basically become the former option— that is, as it is described today, “to coordinate and promote network technology developments and advances in network-based applications and services across the Asia Pacific region…”

Soon after this, however, I moved to the U.S. in early 1997 to become vice president for information technology at IU. I expected my involvement with APAN had come to an end. But in the spring of 1997, the NSF released the HPIIS solicitation that Steve Goldstein had mentioned in Tsukuba. I met with Professor Chon, Professor Goto, Mr. Konishi, and others in Tokyo in June 1997, and we agreed to partner on a joint HPIIS proposal called “TransPAC: A High-Performance Network Connection for Research and

Education between the vBNS and the Asia Pacific Advanced Network (APAN)" on which I was principal investigator.

This was funded by the NSF in 1998 for over $12 million, and substantial additional resources were committed by KDDI and others. TransPAC commenced on July 30, 1998 and was initially a 35Mb connection. Today, it comprises over 120Gbs of connectivity. TransPAC-2 commenced in August 2005, and, as I had taken on additional administrative responsibilities at IU as vice president for research, I asked Jim Williams to take over from me as PI, where he did a superb job until his retirement about seven years ago. Since then, many others have continued this outstanding work, including Steve Wallace, who was involved from the outset, and for the last seven years, Dr. Jennifer Schopf, who now serves as director of International Networks at IU (IN@IU). I think it is fair to say that the credibility that this NSF HPIIS award gave to the whole idea of international connectivity helped significantly with the growth of APAN, and the NSF deserves great credit for this program.

Later in 1998, IU was also awarded the contract to build, manage, and operate the Internet2 network—then called Abilene, and, in 2000, we formally established the Global Network Operations Center—the GlobalNOC—to manage TransPAC, a number of other international connections funded under the HPIIS program, and Abilene. Today, it is the premier NOC in the world for advanced research and education networks. It manages 21 major national and international networks, employs over 130 network engineers and associated technical staff, and has 10 graduate students working on various projects associated with the GlobalNOC.

TransPAC3 was funded in June 2010 along with the America Connects to Europe (ACE) project, also led by IU. ACE had an important role in the discovery of the Higgs boson in 2012 by providing the data bandwidth to move the huge amounts of data from the Large Hadron Collider to the U.S., speeding collaboration and discovery.

Both of these projects were funded again in the next phase of NSF IRNC awards in March 2015, with ACE being replaced by the Networks for European, American, and African Research (NEAAR) project. Both of these projects provided expanded capacity and services to support high-capacity networks and the ever-increasing digital needs of international research collaborations. They also addressed redundancy and resiliency in a globally coordinated fashion, and how high-performance networks could be used more effectively.

IU has recently entered negotiations with the NSF for the follow-on projects to our work in both Europe and Asia as part of the continuing NSF IRNC program, which we hope to be able to announce in the near future. If approved, these projects would allow us to continue our work in Asia, through TransPAC5, and in Europe and Africa with an expansion to the Arctic, through the Networks for European, American, African, and Arctic Research project. We are calling this “NEAAR Next.” Just as TransPAC4 extended the TransPAC3 collaborations to formally include the APAN and TEIN partners, TransPAC5 would build on these and on the significant work of the Asia Pacific Ring collaboration, of which IU was a founding member two years ago. TransPAC5 would extend our multi-lateral agreements to potentially double the capacity and resilience of the Asia Pacific Ring circuits, and offer significantly more focus on supporting extended network experimentation and production research support.

This is just a brief overview of all the numerous global initiatives in high-performance international networking, collectively representing an investment of hundreds of millions of dollars, that can trace their origins back to the enormous success of the original APAN/Indiana University collaboration that established TransPAC. It is no exaggeration to say that without this collaboration and the role that APAN played in it, it is unlikely that global high-performance research and education networking would be as advanced or as inclusive as it is today.

APAN has gone on to become a very important and energetic catalyst and coordinator for the development of regional clusters of national research and education network interconnections in the Asian Pacific region, and for the interconnection of the region to other parts of the world. APAN now includes 17 primary members, with a number of additional primary members to potentially join the organization soon, and with a number of additional affiliate-, associate-, and liaison-members, and a dozen MOU-members.

It has also supported initiatives in fields that include bioinformatics, medical informatics, distributed computing, tele-manufacturing, remote robotic control, digital libraries, wide area parallel computing, astronomy, and high-energy physics. There will be, I know, many talks at this conference about all that APAN has achieved and is achieving.

Finally, let me say as one of the original founders, that APAN and its leaders are to be congratulated for all of their outstanding accomplishments over nearly 25 years. Internationally-based collaboration in science, research, and education in the Asia Pacific region owe you a great debt.

3. WHAT UNIVERSITIES EXPECT FROM HIGH-PERFORMANCE NETWORKS

Now, let me return to a point I mentioned earlier. From a long-term point of view, universities have three fundamental missions:

- the creation of knowledge (research & innovation),
- the dissemination of knowledge (education & learning),

and

- the preservation of knowledge (information repositories).

So where do global high-performance networks fit in here? I would argue that they are fundamental in every one of these areas.

3.1 The Creation of Knowledge

Research has become completely international in character. In the academy, scholarship and research in just about every discipline from anthropology to zoology is truly international—a process hugely accelerated by the Internet. There is, in general, no such thing as American anthropology or Chinese zoology—just anthropology and zoology (though there may be contending schools of theory and analysis within these disciplines).

This scholarship and research takes place within a global research or scholarly environment where it is, in general, facts and reason that determine progress, not national origin. Hence, the quality of the programs and research at universities is determined by the quality of the faculty and students who contribute to them, and they can come from anywhere in the world. And fundamental to research, especially in the sciences, is collaboration, whether it be co-authors on opposite sides of the world, or a group of thousands from dozens of countries working with major experimental facilities located around the globe.

It is also no exaggeration to say that this research and science is becoming almost totally digital. Data is being generated, collected, processed, analyzed, visualized, and stored in digital form. Simulations and modeling are being carried out completely digitally. And the historical and contemporary archives of science, certainly the main material, have been converted fully into digital form.

In this context then, it is the high-performance networks that enable this massive global scientific enterprise to move forward and make progress by ensuring that the speed with which huge volumes of digital data can be moved between researchers and institutions is not a limiting factor to the boldest, most adventurous, and advanced scientific experiments and collaborations that can lead to discoveries and scientific breakthroughs with sometimes profound implications for humanity.

The most pressing challenges we face are also international in character. When it comes to climate change, for example, no one center or group can collect all the data needed for climate models and predictions. International collaborations are an absolute necessity. The same is true for the COVID-19 global pandemic. At IU, our International Networks team has been focusing over the last several months on trying to improve performance from COVID-related data sources in Asia and Europe for collaborators in the United States. Jennifer Schopf, the director of International Networks at IU, will speak more about this work later this week.

3.2 The Dissemination of Knowledge

Just as research has become completely international in character, it is true in many ways of education, with many degrees now requiring some international component, with the rise of global collaborative courseware platforms, with instruction becoming multilateral and virtual—and with all of this fueled, in part, by ubiquitous very high-quality video conferencing and telepresence technologies. As with research, students can also require rapid access to on-line scientific instruments and data that are available to them. As well, until the pandemic at least, we continued to see growth in the number of students in the United States studying overseas, and this is true of many other countries. Nearly a third of the students on our Bloomington campus have studied abroad by the time they graduate. Such international literacy as part of a student's degree has become a very desirable qualification for employers in many fields.

In support of the international aspects of education, many universities have also established campuses and facilities all over the world and they rely on high-quality video conferencing technologies and other advanced communications technologies to continue to closely interact with their home campuses. This becomes even more critical when travel is dramatically curtailed as it is at the moment with the pandemic.

3.3 The Preservation of Knowledge

Of the three missions of universities, the preservation of knowledge is the one that is not as glamorous as the first two, but it is equally important. For the great collections that many of our research universities have are a vital part of the cultural, scientific, and historical records of the human race, and can be among an institution's most precious resources. Increasingly though, the key to the preservation of these resources is through their digitization to allow both their long-term digital preservation or where the long-term survival of the original is at serious risk. One only has to recall the recent tragedies and incalculable losses caused by devastating fires at Universal Studios, the National Museum of Brazil, and Notre Dame Cathedral in Paris.

At Indiana University, we have launched a major effort in recent years to systematically survey and document all of our major collections, and then to build on our early work in digitizing books and documents, to digitize our time-based media objects. This will involve digitizing over 300,000 audio and video recordings, many in formats no longer supported, and around 25,000 films from our world-class film collection. This project is called the Media Digitization and Preservation Initiative, and will be complete next year. We are now beginning to tackle what we think of as the "final frontier" in digital preservation—the comprehensive digitization of 3-dimensional objects. Here, we have a number of flagship projects, including the digitization of large segments of our paleontological collections. We have also carried out a major project with the famous Uffizi Museum in Florence, Italy, to digitize all of the pieces in their Greco-Roman collection as a precursor to our own efforts in this area.

Finally, of course, are the truly massive amounts of "born-digital" data produced by an enormous range of systems and experiments from the smallest lab to the largest particle accelerators, telescopes, or sensor webs. All of this data needs to be seen as part of the overall digital collections and holdings of an institution.

So, in addition to our goal of preserving all these digital objects and digital data, is our goal of providing wide-spread access to them for students, researchers, scholars, and in many cases the general public, who want access to this material. And this is where high-performance networks come in, enabling fast access in a routine manner to massive amounts of digitized data represented by digitized preserved objects in collections.

4. THE FUTURE OF NETWORKED RESEARCH AND EDUCATION

So, given all of this, it is essential that faculty, students, and staff have available to them the infrastructure that will allow them to work at the most sophisticated level and to collaborate effectively with colleagues world-wide. It is impossible to imagine a university or college of any size operating effectively today without it. And during this time when researchers and students are limited in their ability to travel internationally due to the pandemic, I would argue that such an infrastructure is more important than ever.

Conversely, it should by now be unacceptable for this infrastructure to be run at anything but at the highest production standards, as we have come to expect of our national research and education networks, and in the best engineered and best operated way.

However, global networking, that is, the global interconnection of NRENs, has been accomplished overall in an ad hoc way in spite of the best efforts of organizations like APAN and their local successes. There is, on the one hand, over-provisioning in some places caused by local issues of reliability. And until recently, the global interconnecting fabric was, overall, not architected for redundancy or reliability, though over the last five years, improving this situation has been a major focus of the international networking community, especially within the APAN community. The Asia Pacific Ring collaboration between the United States and Asia is a major effort in this regard, just as the Advanced

Network Architecture (ANA) collaboration between the United States and the European Union is for trans-Atlantic collaboration.

On the other hand, there is a lack of connectivity to some parts of the world due to insufficient government funding and unreasonable pricing by commercial entities. As one recent NSF publication states, “Not all communities are connected sufficiently, and those communities often lack the training, relationships, and resources to complete the end-to-end linkages required for full participation in science in the 21st century scientific R&E enterprise.”[2]

A major challenge at the present in the APAN region is in connecting and better supporting the growing economies in countries in the ASEAN region like Vietnam, the Philippines, Indonesia, Malaysia, and more adequately connecting even India and Pakistan to the U.S. and elsewhere. Similar challenges remain for Africa and Latin America.

So, while the dedicated efforts of scores of leaders of national research and education networks have created a global fabric that has met the “first generation” needs of research and education, its new “second generation” demands will require a significantly more systematic and intentional approach to the architecture of global infrastructure—an infrastructure that will need to provide a consistent and seamless advanced set of services, born from a fully integrated set of components. Surmounting these challenges will require innovative approaches to the way these efforts are organized and funded. Achieving these goals will also require the involvement of the key higher education and research organizations in the relevant countries, the key funding agencies in these countries such as the NSF and NIH in the U.S., and other key government and, where possible, multilateral organizations, for international research

[2] National Science Foundation, International Research and education Network Connections (IRNC) Base Program Solicitation, 2020, NSF 20-535, Accessed July 20, 2020, URL: https://www.nsf.gov/pubs/2020/nsf20535/nsf20535.htm.

and education networks are now part of the fundamental fabric not only of research and education within countries, but as part of the global fabric of research and education.

5. LESSONS FROM THE COVID-19 GLOBAL PANDEMIC

Let me conclude with some comments about the coronavirus global pandemic. This pandemic, of course, represents the gravest public health challenge that the world has faced in a century. Here in the U.S., we continue to see the number of cases increase. In the spring, most states instituted some form of “lockdown,” but many states eased their lockdowns too soon, without mask requirements or adequate attention to physical distancing. The pandemic continues to shake the foundations of universities the world over. The rapid pace with which it has developed has been staggering. The unpredictability of events has made an orderly pattern of planning and response impossible. Revenues have plummeted, teaching and research have been disrupted, and international programs have been severely affected.

At IU, like most other American universities, we sent our students home in March and transitioned to all virtual instruction after a two-week break. This was a massive undertaking in an institution with, as I previously mentioned, nearly 100,000 students and 40,000 faculty and staff. However, it went remarkably smoothly with few issues of any kind, even though it was an enormous and unprecedented undertaking.

We became an institution where Zoom and other videoconferencing and collaboration tools became the dominant tools for instruction. To date, for example, we have seen 240 million Zoom minutes – about 450 Zoom years, over 1.1 million Zoom meetings involving over 5.3 million participants, and we’ve seen nearly 15 million Canvas sessions. Many of our researchers, too, had to transition to these and other IT tools to continue their research. This resulted in a huge increase in the load on our campus, local, and regional networks. But they handled them with ease, as did the national networks and the global networking fabric.

Indiana University has invested heavily in its IT infrastructure, including extensively in its high-performance networks, over more than 20 years, and never has the value of this investment been better demonstrated or vindicated than over the last five months of primarily virtual operations. And never has the value of the investments that have been made in national research networks in the U.S. and other countries, like those in APAN, been better demonstrated as the global virtual research and education enterprise barely missed a beat and became the principal vehicle for scientific and research collaboration.

Extremely important to us as well were the investments we have made over 10 years in on-line education through our highly ranked IU Online program where we had many years of success upon which we were able to draw. We transitioned to all on-line instruction against a background where over a third of our students had been doing at least one on-line course per year. It simply cannot be overstated how important this investment was in preparing IU to meet the challenges of the COVID-19 pandemic.

We plan to begin our Fall semester soon with a mixture of in-person and on-line instruction, with strict medical and public health safety measures in place. But, if we face further difficulties in the months ahead, we know we can return to all on-line instruction with great confidence. We will, though, have to continue to operate in this climate of radical uncertainty until effective vaccines are developed or therapeutics that can dramatically reduce mortality.

6. CONCLUSION

Let me conclude by saying what an honor it is to have been invited to give this keynote address to an organization I am very proud to have played a small role in founding nearly 25 years ago. The Asia Pacific faces many challenges in the years ahead, but the spirit of cooperation and collaboration that APAN has so successfully demonstrated since its founding is an excellent example of a positive way ahead. My congratulations to all my old friends who are still among the stalwarts of APAN, and to all the many others who have made it such a success over the years. Finally, congratulations to all at APAN on your 50th conference.

(Professor Shigeki Goto will then lead a 15-20 minute Q&A session.)