ACM SIGACT News Distributed Computing Column 22

Sergio Rajsbaum*

Abstract

The Distributed Computing Column covers the theory of systems that are composed of a number of interacting computing elements. These include problems of communication and networking, databases, distributed shared memory, multiprocessor architectures, operating systems, verification, Internet, and the Web. This issue consists of:

- "Distributed Computing in India," by Sukumar Ghosh.

Many thanks to Sukumar for his contribution to this issue.

Notice that ICDCN `www.icdcn.org` (International Conference on Distributed Computing and Networks), formerly known as IWDC (Intl. Workshop on Distributed Computing), held in India, is trying to attract more PODC/DISC like papers this year. The paper Submission Deadline is May 15, 2006.

Request for Collaborations: Please send me any suggestions for material I should be including in this column, including news and communications, open problems, and authors willing to write a guest column or to review an event related to theory of distributed computing.

Distributed Computing in India

Sukumar Ghosh [1]

1 Introduction

The 90's have seen India emerge as a major player in computing. There were early promises and accomplishments: for example, back in 1957 a fully functional computer TIFRAC was built at the Tata Institute of Fundamental Research, Mumbai, followed by another landmark design ISIJU (in 1961) for the first transistorized digital computer in India, a project jointly taken up by Indian Statistical Institute and Jadavpur University in Kolkata. Qualified manpower was there, and indigenous technologies were slowly blossoming – yet nothing significant happened to make a visible difference. As far as distributed computing was there, telecommunication was at a medieval age, and no significant activity was noticeable in distributed computing – neither in theory nor in practice.

The main thrust in the development of computing and communication technology came from the government sector. National Informatics Center (NIC) was established to act as a focal point for the development

*Instituto de Matemáticas, UNAM. Ciudad Universitaria, Mexico City, D.F. 04510. `rajsbaum@math.unam.mx`.

[1]ghosh@cs.uiowa.edu. Department of Computer Science. 201P Maclean Hall. The University of Iowa. Iowa City, IA 52242-1419, USA.

of methodologies for designing and implementing national information systems and data management techniques. Around 1988, a high-speed network called NICNET was launched (with assistance from UNDP (United Nations Development Program)) to serve as the National info-Highway. NICNET, an incremental overlay over the existing network, was setup with broadband VSAT technology, and since its inception, it has been continuously upgraded using the state-of-the-art technologies [6]. NICNET is now found in all Federal Government departments, and covers the sectors of agriculture, animal health and production, water resources, fisheries, cooperatives, forestry, natural resources, weather data and rural development etc. ERNET, conceived in 1986 and funded by UNDP was the education and research network. It offered email services and access to public domain information. Prior to these, Computer Maintenance Corporation (CMC), an erstwhile government enterprise, set up IndoNet, India's first data network, based on IBM's SNA network. A value-added service provider, IndoNet dealt in optical mark recognition and intelligent character recognition based forms processing services, document management services, file transfer and data management services, etc. IndoNet became operational since 1986.

The capabilities were growing during the 80's. In August 1984, CDOT (Center for Development of Telematics) was established. It used indigenous technology to develop rugged electronic switches that would work without the need for air conditioning, and delivered 32,900 rural exchanges to connect the people of rural India, and telecommunication came of age [1]. Faced with prolonged technology export restrictions by the United States, particularly relating to the Cray supercomputer, another Government organization CDAC (Center for Development of Advanced Computing, founded in 1988) rolled up its sleeves to build the first supercomputer PARAM-8000 in 3 years, and produced the first machine in 1991. The latest machine PARAM 10000 claims to execute instruction at the rate of 100 gigaflops per second. CDAC catapulted India into the realm of supercomputing. Although the cost-effectiveness of these designs is unclear to outsiders, the development of indigenous capability was clearly evident. A comprehensive report by Professor Krithi Ramamritham [5] summarizes the state of computer science research in India till the mid 90's.

The 90's marked a revolution. The commercial web browsers stormed into the scene. Communication and collaboration increased by an order of magnitude. With more bandwidth being available and the cost of the hardware coming down, suddenly the whole world became an even playing field for everyone, including India. Research became more meaningful since they could be readily identified with applications. A second impetus to research came from the emergence of search engines: first Alta Vista and then Google that changed the way we did research. In the past, a major impediment to research was the availability of technical papers. That barrier largely disappeared. Since the late 90's, PC's became available at affordable prices (read: affordable in educational institutions) and the available bandwidth was increasing by leaps and bounds - suddenly the world came to everyone's fingertips.

2 Distributed Computing Research

The Indian researchers have been more active in the networking area. Indigenous technologies to connect the inhabitants of rural India have taken precedence over theoretical issues related to studying systems as a network of processes. This is quite appropriate since research and development funding so far came primarily from Government sources or through Government initiatives. Media Lab Asia (MLA), based in Mumbai, has been setting up wireless networks to take Internet and voice connectivity to India's rural masses. Originally a consortium between Massachusetts Institute of Technology (MIT)'s Media Laboratory and the Indian government (and now owned by the Government of India), MLA has focused on developing and deploying technology solutions appropriate to bridging the digital divide in developing economies. In MLA's research hub at the Indian Institute of Technology (IIT) Kanpur, Professors Pravin Bhagwat and Dheeraj Sanghi launched the prestigious Digital Gangetic Plains project, where they explored 802.11-based rural connectivity. Starting with four villages near Kanpur, they created an "information corridor" between

the cities Kanpur and Lucknow in North India, covering about 25 villages along the route.

Ashok Jhunjhunwala, a Professor of IIT Madras, has been working on a novel technology to connect the rural masses of the vast country: he was planning to use the national railway's network, the phone cable that connects all railway stations as the backbone, and then deploy wireless technology to establish connectivity to the user's house. Since the railway network in India is vast, this will provide a cheap solution to the connectivity problem (source: Fortune magazine, 9 October 2000). Jhunjhunwala's TeNet group works on the vision of providing world-class technology at an affordable price.

Of late, research and development has been attracting private funding from major sources. In 1998, IBM started its Research Lab in New Delhi focusing on areas critical to expanding India's technology infrastructure. Its research initiatives included electronic commerce, e-governance, computational biology, middleware for new software business models, and technologies for human computer interaction. HP Labs India established in February 2002, targeted to enhance the communication infrastructure for small towns and rural areas, improve Internet access, and make information technology available to those who use Indian regional languages. Intel Research created a hub in Bangalore, India's silicon valley, for shaping the next generation computing platforms. In January 2005, Microsoft Research India started its lab in Bangalore – it was their third lab outside US. Last year, Microsoft India funded several basic and applied research projects [2] some of which are more relevant to distributed computing and networking. These include

- (BITS Pilani) Wireless sensor-based water-resource management network for monitoring and control of irrigation
- (IIT, Bombay) Sensor networks for landslide disaster detection
- (IIT, Delhi) Wireless village network for asynchronous communication
- (IIT, Madras) Ad hoc wireless network for village connectivity

Research in ad-hoc wireless and sensor networks is growing at a rapid rate. The Indian Institute of Science (IISc) Bangalore is a leading graduate school with a wide variety of research agenda. Here, Professor Anurag Kumar of the ECE department conducts research on wireless networks in general, including self-organization and self-optimization strategies in sensor networks, and his research gained international visibility. Another institution that is doing significant research in mobile ad-hoc and sensor network is the Indian Institute of Management, Calcutta, Here, Professors Somprakash Bandyopadhyay and Debasish Saha lead an active research group in ad-hoc mobile networks and pervasive computing. Professor Bandyopadhyay maintains ongoing collaborations with a Japanese research group, and his research has received funding from Japanese funding agencies.

At IIT Madras, Professor C. Siva Ram Murthy has a strong presence in Mobile Ad-Hoc and Wireless Networks with numerous publications in International Journals and Conferences covering these areas. Professor C. Pandu Rangan is active in the area of cryptography. His goal is to design secure channels and protocols for multiparty computation, and some of his results have appeared in leading international conferences like PODC, DISC, INDOCRYPT (the Indian version of Crypto). A secondary focus of his research is distributed algorithms aimed at solving graph theoretic problems.

Professor R. Ramanujam of the Institute of Mathematical Sciences (IMSc), Chennai, leads a strong theory group in distributed computing. His research interests are in temporal logics and verification, various partial order models of concurrency and security protocols. Another hub for theoretical research is the Tata Institute of Fundamental Research (TIFR), Mumbai, where Professor RK Shyamasundar leads a group conducting research on specification, verification, model checking and cryptographic protocols.

Professor Krithi Ramamritham of IIT Bombay (formerly at University of Massachusetts, Amherst) made visible research contributions in real-time systems and content distribution networks. After moving to

IIT Bombay, he has been applying concepts from these areas to solve problems in mobile computing, e-commerce, intelligent Internet applications, and the Web. He serves in several editorial boards of journals including IEEE Transactions on Mobile Computing, World Wide Web: Internet and Web Information Systems, and IEEE Internet Computing. Professor Sridhar Iyer of IIT Bombay conducts research on wireless and mobile computing, RFID networks, as well as verification of protocols. Other noteworthy research from IIT Bombay is from Professor Dhamdhere, whose primary interests are in compilers – he has a brief presence in the classical areas of deadlock and termination detection, and more recently, he published papers on self-stabilizing algorithms.

At IIT Kanpur, Professor Pankaj Jalote made past contribution in the area of software fault tolerance for distributed systems. Professor Dheeraj Sanghi's interests are in various networking protocols and load balancing algorithms for web servers. He is also a leader in the Media Lab Asia's rural connectivity project. Professor Bhaskaran Raman is active in wireless networking with a clear presence in some of the top international conferences on networking. Also, Professor Ratan K. Ghosh works in the area of wireless and mobile computing. Professor Manindra Agrawal, whose primary research interest is in complexity theory, created a sensation in the theory community by writing the groundbreaking paper "PRIMES is in P" [4] together with two of his undergraduate students. The result as such has no apparent connection with distributed computing, but it has the potential to make an indirect impact on research in security.

At IIT Kharagpur, Professor Arobinda Gupta conducts research in fault tolerance and adaptive distributed systems. He has several strong publications in the area of self-stabilization, where he made some fundamental contributions on the fault-containment problem. At the newly founded IIT Guwahati, Professor Gautam Barua (who is also the Director of the institute) is currently focusing on security in ad-hoc networks. Also active are Professors G. Sajith (distributed algorithm design and graph theory), Sukumar Nandi and S.V. Rao (wireless, mobile ad-hoc networks).

At the Indian Statistical Institute, Kolkata, Professor Bhabani Sinha heads a research group with diverse interests. Two areas that are of significance in distributed computing are interconnection network and cellular networking. Professor Krishnendu Mukhopadhyaya has been publishing on checkpointing and recovery protocols, and localization problem on sensor grids. At Jadavpur University Kolkata, Professor Pradip K. Das heads the Center for Mobile Computing that promotes basic as well as interdisciplinary research in Mobile Computing (`www.cmccju.org`). He also leads a group that investigates real-time systems, clock synchronization problems and various protocols for checkpointing and rollback recovery. Professor Chandan Mazumdar of the Center for Distributed Computing successfully carried out several projects in fault tolerant and secure distributed computing for Defense Labs and the Indian Space Research Organization (ISRO). These projects weigh more heavily in the developmental component than fundamental research, but certainly play an important role in enhancing indigenous technical expertise in crucial areas. Professor Mazumdar's current interests are in the areas of security, fault tolerance and survivability. Together with Professor Bandyopadhyay of IIM Kolkata, he is also instrumental in developing a rapidly deployable disaster management information network infrastructure using communicating mobile devices, and is running pilot projects in the mangrove islands of Sunderbans in the Bay of Bengal.

3 Conferences

Two conferences in distributed computing are regularly organized in India. The International Workshop on Distributed Computing (IWDC) is running for the past seven years, and recently it has been upgraded to a conference with the new name: International Conference in Distributed Computing and Networking (ICDCN). A much more recent addition is the International Conference on Distributed Computing and Internet Technology (ICDCIT) in Bhubaneswar. On a broader focus, International Conference of High Performance Computing (HiPC) is being organized since 1994. It addresses some topics in distributed

computing, which includes distributed algorithms, grid & P2P networks, and wireless & mobile computing. Foundation of Software Technology and Theoretical Computer Science (FSTTCS) has been running for many years, and it is well recognized as a premier theory conference all over the world. A small fraction of FSTTCS papers addresses theoretical issues of distributed computing. Finally, since the early 90's Advanced Computing Society has been organizing their annual conference ADCOMP. Although it covers a few topics in distributed computing, the conference has a much more general focus.

4 Trends, Funding etc.

In the context of distributed computing, the major emphasis in India seems to be in the area of networking with interests ranging from enterprise networks to mobile and wireless ad-hoc networks and sensor networks. The skilled workforce has absorbed the newer technologies at a rapid pace. The recent growth in bandwidth, connectivity, and deregulation together with the strong programming skills created an upbeat atmosphere in application development involving computing and communication. Fundamental research in distributed computing has just started taking shape. Several factors are fueling this, First, stung by the tough competition and recent funding blues in the US [7], many computer professionals (as well as some Ph.D.s graduating from US universities) are returning to India, and those in distributed computing are continuing their research in Indian institutions. Several non-resident researchers living abroad regularly visit Indian institutions, and find the intellectual environment at the top institutions quite stimulating. Short-term collaborations during these visits spawn and invigorate research. Third, the funding scenario in computing is less competitive in India: most researchers with some bright ideas are able to secure funding either from one of several government agencies (DST, DIT, ISRO, Defense), Media Lab Asia or from the Indian offices of the multinational companies like Microsoft or IBM or Intel without much problem. In addition, joint proposals submitted to NSF and DST (Department of Science and Technology, Government of India) have received favorable consideration so far. This is in sharp contrast with the current highly competitive funding situation in the US. Fourth, the international conferences in distributed computing are well attended by local graduate students and researchers – these generate enough interest and awareness in new research directions. Fifth, salaries in computing and information technology have witnessed significant growth: enough to attract good researchers from abroad or dissuade bright graduates from leaving the country. IIT graduate are still much sought after in the graduate programs of top US institutions like MIT, Stanford, or University of California Berkeley, but the exodus is slowing down.

5 Conclusion

In a country where 70% of the people live in the villages and a substantial fraction of the population are not literate, how relevant is distributed computing research? For the rural sector, the emphasis has been (justifiably) in improving connectivity before distributed computing research makes any sense. Accordingly many villages have set up cell phone or laptop and Internet kiosks where people can walk in to use the facilities. In some areas, mail carriers carry cell phones using which villagers can make phone calls for a nominal charge. One may recall that in 1999, Indian scientist Swami Manohar invented the Simputer [3], a low cost handheld computer, to bridge the digital divide. Linux-based Simputer has multi-lingual capabilities, text-to-speech software, and simple handwriting recognition software. Simputer maker Picopeta introduced the first retail version last April with a price tag around $250, but the sales fell far short of the expected figures. One may compare this with a similar project for a $100 laptop by the MIT Media Lab for the Delaware-based non-profit organization One Laptop Per Child (OLPC).

Indian research and development is primarily feeding the multinationals, and other than improving con-

nectivity and developing interfaces for Indian languages, not many applications are directly channeled to the Indian panorama. The theory front has a solid base, and with new researchers joining the universities, IITs and various research organizations, a steady growth is expected. On the experimental side, one has to wait and see how the present euphoria shepherds distributed computing research to bear fruit relevant to the local context. The realization that there is a tremendous scope for developing applications involving mobile devices, sensors and RFIDs has already dawned.

6 Acknowledgements and Disclaimer

The report is not comprehensive, but best effort has been made to collect relevant data through emails, public media, and personal contacts. Any omission or error is inadvertent. The author is thankful to Arobinda Gupta, Anurag Kumar, Sajal Das, Chandan Mazumdar, RK Shyamasundar, Somprakash Bandyopadhyay and Sergio Rajsbaum for many comments, inputs and useful pointers.

References

[1] Dataquest http://www.dqindia.com/content/20years/102122311.asp

[2] Microsoft Research Launches Operations in India (Microsoft News) http://www.microsoft.com/presspass/press/2005/jan05/01-12MSRIndiaLaunchPR.mspx

[3] "Indian handheld to tackle digital divide." (July 18, 2001). BBC News: http://news.bbc.co.uk/1/hi/sci/tech/1442000.stm

[4] Manindra Agrawal, Neeraj Kayal, Nitin Saxena, "PRIMES is in P." Annals of Mathematics 160(2): 781-793 (2004).

[5] Krithi Ramamritham, "Computer Science Research in India," http://www.fas.org/nuke/guide/india/agency/krithi1.html November 1995.

[6] Homepage of the Satellite Communication Group of NIC: http://satcom.nic.in/

[7] Rick Merritt: Computer Science R & D goes begging for funds: http://lazowska.cs.washington.edu/eetimes2.htm