

New applications of broadband networks: A vision based on the state of the art in networking

B. Plattner

*Computer Engineering and Networks Laboratory, ETH Zurich
ETH-Zentrum, CH-8092 Zurich*

Phone +41 1 632 7000, Fax ext. 1035, plattner@tik.ee.ethz.ch

Abstract

In this paper I discuss potential research and development directions that may prove relevant within the next ten years. I start with an observation documented by the history of research and product development in information technology, which indicates that it will take approximately ten years of product and market development efforts to transform a well-working laboratory prototype into a successful product. Thus we can estimate what the information technology market will offer ten years from now. Similarly, we can make a prediction about the topics that will be relevant in applied research by reviewing today's long-term research areas.

Keywords

new applications, vision, prediction, broadband, information technology

1 INTRODUCTION

I have been invited and asked to develop my vision of networking in general and, specifically, of broadband communications. While I agree that researchers need to have visions of the not-too-near future - definitely beyond the rather short-term visions that shareholder value oriented organizations tend to have - it is not implied that this requires prophetic capabilities. Visions, in our profession, should be based on evidence - but what evidence should we consider in our fast-living field, where one (Internet) year is said to be worth seven normal years?

To be on the safe side, I will try to generate a vision of the networking scene as it will present itself in the market ten years from now. To this end, let me start with the observation that most of today's established technology was well understood in the research lab roughly ten years before it hit the market, i.e. there appears to be a time constant of ten years for a product to mature from its invention as a basic technology. There are various examples that illustrate this observation:

- Ethernet technology was developed by researchers at the Xerox PARC in the early seventies (see Metcalfe and Boggs, 1976, and the on-line replica of US Patent No. 4063220); however LAN products were available in the market only after 1980, and became an affordable mass product even later.
- The timeline of the Internet is similar. Its key concepts and initial implementations were developed by Cerf and Kahn in 1973 (see Cerf and Kahn, 1974, and Cerf, 1993), and were deployed in a rather restricted academic and military environment in the following years; only between 1985 and 1990, IP routers started to be available in the market from companies like Cisco, 3Com and Wellfleet.
- The history of graphical user interfaces as we know them today again is alike: Developed in the early seventies (again at Xerox PARC), they started to be available for normal users with the arrival of the Apple Macintosh, which was launched in 1984 (see <http://www.apple-history.pair.com/128k.html>).

It may be argued that all these examples are similar as they pertain to basic technologies developed in the early seventies. However, the deployment of ATM technology shows a similar pattern: Early papers on ATM, a technology initially called fast packet switching or asynchronous time-division multiplexing, were published as early as 1984 (see Thomas, Coudreuse and Serval, 1984, for a discussion of an experimental system, or Turner, 1986, for an overview of early research), but reasonably priced equipment was not available in the market before 1995.

It is obvious that the interval of ten years is not just needed for readying a technology for the market, but also for the creation of the market itself, and for refining the production processes which will enable the industry to deliver products at prices that are acceptable to the consumers.

2 THE STATE OF THE ART IN NETWORKING

Following the observation given in the introduction, let's consider the technologies under study in the research labs at this time, and we will know what we will use ten years from now. In the sequel, I will also identify steps to be taken in the near future to further each of the technologies considered.

Broadband networking (Gbit/s backbone and Mbit/s access)

In the research labs, the basic technologies for broadband networking have come of age. It will be the challenge of the coming years to deploy Gbit/s technology to provide a universally connected infrastructure at affordable prices. An crucial issue will be to implement broadband access networks (using traditional copper, optical or wireless media) that will allow users to take advantage of the advanced applications possible with broadband networks.

Security technology

Security technology in the form of basic cryptographic mechanisms has been avail-

able for some time; however, it was never really deployed due to several reasons: Political measures (e.g. export or usage control applied by various governments) were - and are still - an obstacle preventing their general use. In addition, the market pressure has not been strong enough to create the infrastructure needed for key distribution (e.g., trusted certification authorities). The situation is shifting rapidly with the introduction of electronic commerce, which calls for a broad use of security technology for confidential and committed exchange of information. Besides the deployment of a suitable trust infrastructure, the development of workable and scalable digital payment systems (also for micropayments) will be major challenges. We are still far from being able to do digitally all we can do with traditional forms of money and value representation.

Distributed multimedia applications and integrated services networks

These topics have been very hot ones in the last few years, but despite great efforts the capabilities of computer-based multimedia systems and integrated services networks needed to deploy distributed multimedia are still far from satisfactory. Audio and video quality are often poor, delays due to encoding, transmission and decoding are too long, and signalling systems for providing QoS in networks have become much more complex than anticipated (this applies both to ATM signalling and resource reservation protocols in the Internet). I think that the breakthrough is still ahead of us, and new approaches are dearly needed.

Assuming that these problems will eventually be solved, a new problem awaits a solution (especially in the scope of the Internet): A network providing different service classes will only be workable if their usage is charged to individual users or to groups of users. However, the importance of user-based pricing, charging, accounting and billing in open and cooperatively operated networks have only recently started to draw upon the attention of the research community.

Optical networks

Unlocking the power of optical transmission for communications has been one of the major goals of many research projects. It may be anticipated that the deployment of Wavelength Division Multiplexing (WDM) technology and advances in Code Division Multiple Access (CDMA) will provide the basis for many services requiring broadband access and transmission. Nevertheless, functionally replacing current electronic switching technology with all-optical switching lies far ahead (see section 3).

Mobility

Mobile systems have been one of the fastest growing markets in the past, especially for voice communication. It is obvious that the fascination of mobility will also apply to all kinds of computer based communication, leading to facilities as we know them from Startrek and other science fiction movies. Various problems still need to be solved, however, such as broadband wireless multiaccess transmission, QoS management for multimedia applications, and providing an architecture for a mixed

mobile/terrestrial infrastructure offering a seamless and reliable service to the travelling user. Besides humans, also machines (processes) will be mobile communication partners, and parts of the infrastructure will be mobile as well (just imagine airborne users accessing the Internet via an airborne subnetwork).

Flexible and easy service creation

The explosive growth of distributed multimedia applications, such as video conferencing, video-on-demand, and collaborative distributed environments, has led to a variety of networked multimedia services. To support this growth effectively, both the network and the end systems require flexible service-enabling platforms, also referred to as middleware. These platforms include the functions of accessing and sharing various types of resources, and they provide generic support for service creation, service delivery and service management. The task of such a platform can be compared to that of an operating system providing support to applications on a computer.

Service enabling platforms are being made possible by the convergence of several technologies, specifically, distributed computing, multimedia systems and telecommunications architectures. The challenge will be to develop and build sufficiently high-utility and high-performance platforms that run on top of advanced networks with fixed as well as mobile/wireless components¹.

Component-based software

As illustrated with developments in the Internet, there is a trend towards a decomposition of full-blown applications into smaller components, which can dynamically be activated and integrated into a given software environment. Java applets and the (marketing) promise of the network computer are just the beginning of a development which may well create a whole new paradigm of software reusability and for the way we create applications. If successful, these new approaches will lead to the dynamic composition of multi-functional applications and pave the way for pay-per-use software. The main incentive to pay for each use would be that this economic model requires a minimum investment by the customer, of course offset by a higher cost for actual usage. Many users may take advantage of such services, or will prefer a mix of pay-per-use and pay-for-ownership models. It is obvious that component-based software can only be implemented with the security services that were previously mentioned.

1. This text is taken from the Call for Papers of an upcoming issue of the *IEEE Journal on Selected Areas in Communications* on the theme "Service Enabling Platforms for Networked Multimedia Systems". The author is one of the guest editors of this issue.

3 NEW APPLICATIONS AND DIRECTIONS FOR FUTURE RESEARCH

The developments above, which are not speculative but based on work currently carried out in research and development laboratories, will be complemented by the ever increasing complexity and performance of semiconductor chips, accompanied by a commensurate decrease in per-function prices. They will enable new applications, such as:

- Multimedia-based, mobile group communication, whenever and wherever we wish or need to have it, with service guarantees (if at all required by the user).
- New forms of entertainment systems, in which the public may actively participate (e.g. with the use of virtual reality approaches). Traditional TV as we know it today may disappear altogether (however, not within the ten years considered here).
- The technological progress will enable applications needed for distance education: Virtual tele-presence will be possible without restriction in quality or quantity, and instruction material and live simulations will readily be available to the student. By consequence, dramatic changes in our systems for higher and continuing education have to be expected, changes that will probably be relevant for most of the readers that are not yet close to retirement. In line with changes in education, the same technology will foster a shift from the traditional workplaces to electronic home workplaces, which may be a preferred way of working for many people active in information-oriented businesses and professions.
- It is apparent that the long-awaited convergence of computers and communications is now happening, triggered by the Internet wave that swept the world in the past few years. We will see that technologies and applications from the areas of database systems and information retrieval will be joined with methods from telecommunications and computer networking. The former disciplines traditionally used architectures with central control, and gradually introduced aspects of distributed systems, while the latter started out with a distributed approach and attempted to achieve reliability and consistency of information, providing to the user the illusion of a centralized system. It is my strong belief that the gradual transformation of the Internet from a computer network infrastructure (providing connectivity and transport of information as a primary service) to a huge information resource is just a beginning. Whenever we realize how difficult it actually is to find the right information fast, using today's Internet search engines, it becomes obvious that we still have a long way to go.
- New approaches to control systems will evolve, likely to be denoted as "remote xy", where "xy" equates to any member of the set {surgery, factory control, control of transport systems, etc.}.
- Finally, electronic commerce will be a booming domain, especially since in the

information society many goods to be traded will be information. Such trade is easily done on networks, and will generate the need for new small or large scale systems for general trade, brokerage, funds management, and clearing, etc.

However, some of the more far-reaching research carried out today will not have its effect within the next ten years. Just consider a few of these areas (without striving for completeness, of course):

All-optical (photonic) switching

While photonic switching has made some progress lately, especially on the component level, system-oriented work is still some years away. However, if available, all-optical networks would turn the wheel of development back in certain respects: Optical networks are analog by nature, tend to be noisy (due to optical amplifiers) and are subject to nonlinearities (see Thylén et al., 1996), as opposed to the digital representation possible today in electronic switching and computing. This means that we will probably see a combination of photonic and electronic switching systems, as soon as photonic switching components will become available. One specific concern, which is due to the noise problem present in optical switching systems, is their limited scalability, which will dictate the establishment of photonic switching domains interconnected with optoelectronic converters.

Active networks

The concept of active networks has fast become a hot research topic. The incentive to make networks active is to enable users or user groups to tailor a general networking infrastructure to the need of their specific applications, e.g. by embedding application-oriented functions into the processing nodes of a network. This is well connected with the topic of (dynamic) service creation discussed above, as it will allow to implement services to be offered by the network dynamically and on demand. However, much research needs to be done in this area, and it is not obvious that all these efforts will prove beneficial, as considerable difficulties have to be overcome. Security will be a primary issue, as well as the efficiency, correctness and reliability of the services to be implanted into an active network.

Quantum computing and quantum cryptography

Quantum computing¹ has the (faint?) potential of liberating us from the sequential nature of computing in which we have been imprisoned since von Neumann's invention (sequential computation is the rule even in the area of massively parallel processing, which just replicates sequential processors by the thousands). Quantum computing, if ever available, would offer mechanisms to solve NP-hard problems in polynomial time, as a quantum computer would be capable of performing an arbitrary number of operations in parallel and link the results in a probabilistic manner.

1. see <http://aerodec.anu.edu.au/~qc/> for an updated bibliography on quantum computing

As a side-effect, quantum computers would render obsolete public-key cryptosystems as we know them today, since these rely on the impossibility to compute the solutions for certain large problems. But, as a bonus, quantum-based systems could provide new methods for the exchange of key material. Quantum cryptographic systems take advantage of Heisenberg's uncertainty principle, according to which measuring a quantum system in general disturbs it and yields incomplete information about its state before the measurement. Monitoring a quantum communication channel therefore causes an a disturbance which cannot be prevented by the eavesdropper, alerting the legitimate users. This yields a cryptographic system that can be used for the distribution of a secret random cryptographic key between two parties initially sharing no secret information. Research on quantum cryptography was pioneered by Wiesner (Wiesner, 1983); Gilles Brassard provides an excellent overview of the subject (see "A Bibliography of Quantum Cryptography", <http://www.iro.umontreal.ca/~crepeau/Biblio-QC.html>).

4 CONCLUSIONS

The vision given in this paper certainly is a forward-looking one. It should be obvious that the state of the art, as sketched in the respective section of this paper, still exposes many unsolved problems and opportunities for research. In the past we have achieved progress in all these areas, but to actually make the promises of the technologies currently under study come true, much research is required, primarily application-oriented research in which systems are designed, implemented and evaluated.

One should not under-estimate, however, the implications of economic, political and legal constraints on the deployment of technology. In parallel with technology deployment, a legal framework fitting the needs of the information society has to be established, not only within countries, but also on an international level.

In my deliberations, I have met some assumptions about the economic and political framework:

- The liberalisation of the telecommunications market which was established in Europe at the very time of writing will eventually lead to a distinct decrease of per-unit price for bandwidth. Broadband applications will not be possible if the unit price for bandwidth were still tied to the price of a simple telephone connection, as has been the case in the past of the monopolistic regime in telecommunications.
- Competition in the IT market will not be hampered by emerging de-facto monopolies, such as the Microsoft-Intel complex and the follow-ups of monopolistic telecommunication carriers which can already be identified after the magic date of the 1st January 1998.

Assuming this, and also anticipating that no global catastrophe will hit us in the near future, I think that the field of broadband networking will remain a very interesting

and active area, in terms of research, development of products and the impact on our society as a whole.

5 REFERENCES

- Cerf, V. and Kahn, R. (1974): A Protocol for Packet Switching Network Interconnection, *IEEE Transactions on Communications*, vol. COM-22, pp. 637-648, May 1974.
- Cerf, V. (1993) How the Internet Came to Be, in "The On-line User's Encyclopedia" by Bernard Aboba, Addison-Wesley, Nov. 1993, ISBN 0-201-62214-9.
- Metcalfe, R.M, Boggs, D.R. (1976) Ethernet: Distributed Packet Switching for Local Computer Networks, *Commun. of the ACM*, vol. 19, pp. 395-404, July 1976.
- US Patent No. 4063220,
http://patent.womplex.ibm.com/details?patent_number=4063220.
- Thomas, A., Coudreuse, J.-P., Servel, M. (1984) Asynchronous time-division techniques: An experimental packet network integrating videocommunication, *Proc. Int. Switching Symp. (ISS)*, Florence, May 1984.
- Thylén, L., Karlsson, G., Nilsson, O. (1996) Switching Technologies for Future Guided Wave Optical Networks: Potentials and Limitation of Photonics and Electronics, *IEEE Communications Mag.*, Vol. 34, No. 2, February 1996.
- Turner, J. (1986) New Directions in Communications (or Which Way to the Information Age?), *Proceedings of the Zurich Seminar on Digital Communication*, pp. 25-32, 1986.
- Wiesner, S. (1983) Conjugate coding, *Sigact News*, vol. 15, no. 1, 1983, pp. 78 - 88; original manuscript written circa 1970.