

THE FUTURE OF HIGH-PERFORMANCE MEDIA NETWORKING

KEVIN GROSS¹

¹ *QSC Audio Products (consultant), Boulder, Colorado, USA*
kevin.gross@avanw.com

High-performance media networking represents an emerging class of network application. The future of these applications is dependent on the trajectory of technology development in the IT industry. Market factors will also play a significant role in determining their success. By examining past development in the network field, the fundamentals of network economics and the requirements of media applications, it is possible to make some predictions about the future of network technology and networked media applications.

INTRODUCTION

High-performance media networking refers to real-time high-quality transport of live digital audio and video over data communications networks. The intent is to replace exiting AV-specific cable plants in commercial and broadcast facilities with a standard data communications network infrastructure. In order to accomplish this, certain performance criteria must be met. The performance of network technology has improved exponentially throughout its existence and is expected to continue to improve along this path. The higher performance opens new opportunities for media networking applications.

1 EVOLUTION OF NETWORKING

The field of networked communications originated as two separate disciplines. Data communications (“datacom”) technology evolved from the simple point-to-point digital systems used to connect early computers to one another, to the meshed networks we use today. Modern telecommunications (“telecom”) technology evolved separately starting from the telephone system. At the same time the computer networks were being commercialized in the 1980s, the telephone system made its transition from analog to digital. The result was two separate digital network disciplines. As the capabilities of each discipline advanced, interconnections between the two were made: Telephone connections were used for computer networking (e.g. voice-band modems, ISDN) and computer networks began to carry voice and multimedia communications (e.g. VoIP, Internet teleconferencing).

Realizing that the voice-oriented telephone technology was not flexible enough to effectively accommodate data communications, telecom made a substantial but eventually unsuccessful effort to unite the fields with the development of a revolutionary digital network technology called ATM.

Meanwhile, datacom networks were quietly and incrementally enhanced to provide higher bandwidth and stable performance for the carriage of real-time voice and media. This has enabled datacom networks, with increasing success, to perform telecom functions. As a result, telecom applications are now being absorbed by new datacom technology.

2 DATA COMMUNICATIONS TODAY

Today it is clear that future advances in digital communications will be rooted in datacom technology, and more specifically, IP networks. The strength of today’s datacom networks is their simple and robust end-to-end operating principle, high-bandwidth capability and ability to adapt rapidly to changing use patterns. Their weakness is in a lack of service guarantees, an approach known as “best effort”.

Most LANs today are built with Ethernet. Ethernet has evolved from a limited system based on unwieldy coaxial cables to a scalable high-bandwidth system based on unshielded-twisted-pair and fiber-optic cables. Network equipment has evolved from passive connectors and wiring to include sophisticated switching and routing equipment capable of inspecting and moderating high-speed traffic.

The Internet Protocol has been in universal use on the Internet since the 1980s. However, through the mid-1990s, until the industry settled on IP, there were numerous network transport protocols in use on LANs. These alternate transports are now largely obsolete and network administrators now generally assume modern protocols to be IP-based. The dominance of IP has blurred the line between the LAN and wider-area networking.

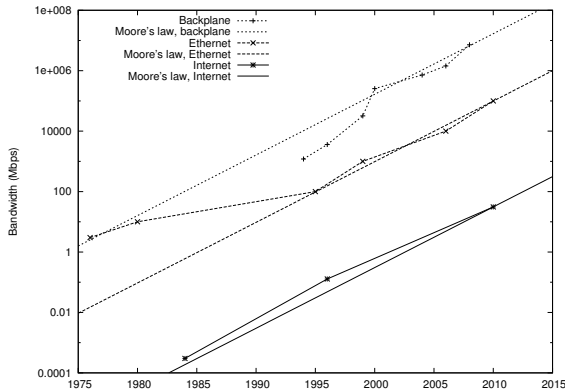


Figure 1: Growth of backplane, Ethernet and end-user Internet access compared to Moore's law. In response to increasing compute power and server density in the datacenter, backplane bandwidth has exceeded Moore's law. The fact that LAN bandwidth to individual workstations is not often a computing bottleneck explains the slower growth of Ethernet at least until it began to be used in the datacenter. End-user Internet access speed has closely tracked Moore's law.

The capacity of computer networks has shown an ability to grow with demand. In pure computing applications, communications demand grows in relation to Moore's law.¹ Network capacity in datacenters matches or exceeds Moore's law.[1] Most network applications are not pure computing applications, however, and the bandwidth of individual Ethernet connections used for general connectivity are not pushed as hard and have historically grown at a slower rate.

3 NETWORK ECONOMICS

The *network effect* is an increase in value of a network connection to each participant with each new participant that joins the network. The telephone network illustrates this well; A telephone is not valuable to you until the people you want to talk to also have telephones. When a system like this grows, it almost magically generates value which fuels further growth, produces positive feedback and further increases growth rate. The network effect is responsible for a number of non-intuitive economic phenomena including the plethora of free services on the Internet and the victory of VHS over Beta.[2 ch.7]

¹ Gordon Moore predicted a doubling of the number of transistors on integrated circuits every 2 years. The relevant metric for computing is not transistor count but overall performance of the electronics which is predicted as a corollary to Moore's law attributed to David House as a doubling every 18 months. This 18 month doubling is used for Moore's law calculations in this paper.

Systems subject to the network effect are characterized by a long early-adoption period followed by explosive growth. This perhaps explains why we have so far seen a fragmented market and relatively low adoption rates through the 15 year history of audio networking. It predicts we're in for a ride in the future as the following anecdotes demonstrate.

- FAX technology was invented in 1842. It was not commercialized until 1925. It wasn't until 1982 that private FAX machines became readily available. Five years later, half of all businesses had FAX.
- Although the first e-mail was sent in 1969, e-mail did not have a significant presence until 1989. It doubled every year thereafter until 1995 when it truly exploded.[2 p.13]

4 INTEROPERABILITY

Interoperability can have a dramatic effect on network economics. By opening divisions between separate networked systems, interoperability immediately increases the effective size and value of a network. Because it merges entities, interoperability reduces the variety of individual technological pieces. This loss of diversity is more than compensated for by an increase in the variety of ways pieces can be assembled to build a system. Increased versatility along with improved economies of scale are manifestations of the network effect. Since users experience increased options at reduced prices, interoperability invariably benefits the consumer. Unless the effort required to achieve interoperability is large, the value imbued by the network effect improves margins and generally benefits industry as a whole. Individual participants may see interoperability as a threat to their existing (locked-in) business. One of the reasons networked markets take a long time to develop is that interoperability can be vetoed by a single influential and uncooperative market participant.

5 MEDIA NETWORKING TODAY

The ability to transmit high-quality real-time media over today's data networks depends on the ability of the network to provide adequate bandwidth and responsiveness. Until recently, there were three approaches to assuring adequate performance.

1. Build a sufficiently capable dedicated network for each media application.
2. Build a shared network with enough capacity to simultaneously accommodate the worst-case traffic patterns of all applications (a practice known as "overprovisioning").
3. Design applications capable of tolerating the variable performance experienced on a shared network.

Option 1 (dedicated network) is used by early media networks such as the original CobraNet implementation, EtherSound and AES50. Option 2 (overprovisioning) is used by current CobraNet and is presently still a viable option for many network designs. Option 3 (application resilience) is used for streaming media over the Internet (e.g. Internet radio, YouTube) but is generally unable to meet the reliability and performance requirements of commercial applications.

With the continuing evolution of datacom technology, a fourth option has recently become available:

4. Quality of service (QoS)

QoS mechanisms mark and prioritize individual packets and traffic streams. These mechanisms allow a shared network to deliver nearly the same performance to critical applications as a dedicated network.

6 NETWORKING IN THE FUTURE

Extrapolating from past experience, the current state of networking and the economic rules governing networked systems, we can make the following predictions.

6.1 Security

Network designers and administrators from the IT profession will have significant influence on the shape of future networks. IT professionals tend to view the world as one network. This view will prevail and the standalone, dedicated audio network will eventually cease to exist. Applications served by dedicated networks today will operate on the one network of tomorrow but will be significantly protected from interfering with one another or unintentionally sharing data in much the same way that applications are protected from each other on modern operating systems. On computers this is accomplished using memory segmentation and protection layers in the processor. On networks, it is accomplished using VLANs, IP subnets and routing, security protocols and access controls.

These security measures are implemented at the edges of the network and will be increasingly visible to network applications and their human creators and caretakers. Standard protocols developed by the IT industry will more easily traverse emerging security provisions. The path of least resistance for applications developers will be to reuse existing protocols and to collaborate with network engineers to produce working systems. This preference for standard protocols will make RTP/IP the dominant transport protocol for live media applications.

6.2 IP

Increased emphasis on security and administration will further IP's position as the dominant protocol suite. Non-IP protocols may be considered a vector for cyber-attacks. Out of concern for security, network adminis-

trators will block applications that do not use proven IP protocols. Non-IP protocols will have no general assurance of connectivity - they will simply not work on some networks. Future network development will be nearly exclusively IP based.

The lower-layer protocols marginalized by these changes can be upgraded to use IP with a moderate increase in overhead. IP protocols are readily carried on lower layer networks but low layer protocols will go nowhere on newer secured and converged networks. By using higher layer protocols, applications gain flexibility while losing little.

IPv6 will be deployed to overcome scalability issues with the current version of Internet Protocol (IPv4). Although IPv6 uses the same engineering principles as its predecessor, it is expected to be a complex and lengthy transition. As IPv6 is deployed and some of the barriers in the Internet created by IPv4 limitations are removed, the distinction between LAN, WAN and Internet will weaken further.

6.3 Synchronization

We have thus far relied on a number of techniques for distribution of time for our media systems. These include point-to-point links, radio and satellites. NTP has been used to transmit time over the Internet since the beginning of the Internet. Clock delivery over networks has recently advanced with the development of IEEE 1588 and further in the IEEE 802.1AS standard used by AVB. Additional means of high-precision clock delivery over wide areas will be developed. As with our other communications endeavours, these are likely to be datacom oriented. Our capacity for world-wide synchronization will improve to the extent that accurate time will be a standard service provided by any network connection.

6.4 Transport




Figure 2: Due to advances in data compression technology, growth of media bandwidth shows no clear pattern and is overall flat in comparison to growth of network bandwidth (Figure 1 and *Moore's law, all media* plot in this figure). Bandwidths and date of introduction of audio formats are shown in lowercase;

Video formats are shown in uppercase: a-digital telephone, b-CD, c-GSM, D-MPEG-1, E-DirectTV (MPEG-2), F-terrestrial digital HDTV (US), G-DVD, h-iTunes music, i-XM radio, j-Skype audio, K-Blu-Ray, L-Skype video, M-Skype high-quality video, n-improved iTunes music, O-Skype HD.

The growth in bandwidth requirements for individual media streams has slowed as encoding technology continues to improve² and as we reach the point of diminishing returns where advanced AV formats meet or exceed the capabilities of human senses. This slowing is partially offset by the fact that the number of channels used in media systems continues to rise. These larger systems require more capable and more extensive distribution infrastructure. Network technology, already geared for deployment on a worldwide scale, clearly exceeds these requirements. Yet, network capabilities and capacity will continue to advance. Along with an increase in bandwidth will come decreased latency and improved performance in other dimensions (e.g. fault tolerance). Due to improved economics and availability, networked media distribution will become viable in a greatly widened swath of applications.

It is worth noting that environments such as this where capability outstrips demand can give rise to disruptive innovation. Although we can expect disruption, we cannot predict its exact nature other than it will serve to close the gap between network capabilities and network application requirements.

² Although storage and wired networking arguably no longer require it, improvements in data compression are driven by mobile and satellite applications where bandwidth is fundamentally limited or expensive.

On the other hand, because networking is largely about interoperability, milder evolutionary changes have a distinct advantage over revolutionary ones. Although possibly advanced in unexpected ways, tomorrow's network will remain recognizable to us.

6.5 Quality of service

“The Holy Grail of computer networking is to design a network that has the flexibility and low cost of the Internet, yet offers the end-to-end quality-of-service guarantees of the telephone network.”[3] There have been multiple attempts to build QoS capabilities in accordance with this goal (e.g. ATM, IntServ, Premium Internet2 service). None have clearly succeeded. Even the scaled-back goals represented by DiffServ have been largely abandoned in wide-area deployments. In addition to the formidable technical difficulties, QoS on a WAN scale involves economic and political factors that further tangle the problem.[4] As a result, QoS progress has been slow in comparison to the progress advancing network technology in other dimensions (e.g. improving bandwidth).

Despite these challenges, QoS functionality in network equipment has stubbornly improved. We now have enterprise and even small-business equipment capable of the high-performance traffic prioritization, shaping and policing required to implement a competent QoS system. In the context of a privately-controlled campus network, a combination of overprovisioning and static configuration of these capabilities by network managers now provides useful QoS functionality. This approach is commonly known as an *engineered network*.

AVB-style QoS with admission control is useful up to a LAN scale. The more scalable DiffServ approach is useful on a campus network. Given a toehold on the LAN and campus and prodded forward by VoIP and IPTV applications, QoS features will continue to improve. Network management software will advance to assist with QoS troubleshooting and tuning. Once a critical number of applications are using QoS, the accompanying mushrooming administrative overhead and unexpected interactions will motivate additional effort to improve our understanding of the general QoS problem. On these medium-scale networks, there is nothing fundamentally stopping us from approaching the Holy Grail of networking. However, this will require concerted effort by academics and then standards bodies. These efforts will undoubtedly produce additional failures along with some advances.

Viability of QoS initiatives on wide area networks including the Internet are hampered by a host of additional complications. These networks operate in a region of diminishing returns where QoS becomes exponentially more expensive as the network scales and linearly less valuable as increasing propagation times

over the distances involved outweigh gains achieved through QoS.

7 MEDIA NETWORKING

The state-of-the-art in media distribution has consistently lagged the state-of-the-art in telecommunications. Following in the path of telephone technology, media distribution has progressed from dedicated analog connections configured through patch bays to a digital circuit-switched solution based on central routers. Extrapolating from the telecom experience, it is clear that media distribution will follow telecom onto the packet-switched datacom networks.

At present, forces preventing media distribution from completing the leap to datacom networks include the additional expense and complexity of interfacing to a network compared to a simple point-to-point connection, a fragmented market, poor economies of scale and current performance limitations of data communications networks in terms of bandwidth, QoS reliability and security.

With increasing network performance, more media applications will get onto networks. With more devices and components and software development environments becoming network oriented, there is a lower barrier to entry for media network applications. Media will be the principal user of QoS services. Networks and media applications will evolve in cooperation to enable advanced capabilities.

8 CONCLUSIONS

Networking will become more central and more singular to everything we do. One highly interconnected, high performance, highly reliable network will be integral to commercial facilities of the future. Security and administrative frameworks will give applications the impression of having a dedicated, high-capacity and secure infrastructure on which to operate.

LANs and campus networks will adopt more of the technology used in WANs. High-performance media streaming applications will need to use IP protocols and techniques to ensure compatibility with the network infrastructure. There is little downside to an IP-centric implementation.

Network performance will continue to improve. However, for general networking over wide areas, best-effort service and the end-to-end principle will prevail. QoS will be deployed within local domains up to campus scale. These networks will be tuned to provide real-time performance for applications that require it. Real-time performance capabilities will mature over the coming years requiring reduced administrative intervention and engineering.

Network economics indicate that adoption of network applications such as high-performance media streaming is lengthy process. Interoperability initiatives are a

potential shortcut to reaching the critical mass where exponential growth begins.

High-performance media networking is ripe for disruptive innovation. While the exact nature of these advances is unknown, they will likely exploit emerging network capabilities in bandwidth, QoS and security.

REFERENCES

- [1] Omar Sultan (25 March 2009). "Networking: Delivering More by Exceeding the Law of Moore." Cisco. Archived from the original on 13 June 2011 at <http://www.webcitation.org/5zPfkCHxA>
- [2] Carl Shapiro; Hal L. Varian (1999). *Information Rules*. Harvard Business School Press. ISBN 9780875848631.
- [3] Srinivasan Keshav (1997). *An engineering approach to computer networking: ATM networks, the internet, and the telephone network*. Addison-Wesley. ISBN 9780201634426.
- [4] "Why Premium IP Service Has Not Deployed (and Probably Never Will)". Internet2 QoS Working Group. 3 May 2002. Archived from the original on 13 June 2011 at <http://www.webcitation.org/5zPijl6Mn>