

Diversinet: Embracing Heterogeneity in Future Network Services

Colin Perkins Allison Mankin
USC Information Sciences Institute
4350 N. Fairfax Drive, Suite 620
Arlington, Virginia, 22203
Email: {csp|mankin}@isi.edu

4 February 2001

Abstract

We consider the problems posed by the increasing heterogeneity of network services and applications, in particular the convergence of data, voice and televisual networks and applications. A possible approach by which this heterogeneity can be tamed and taken advantage is to embrace diversity at the levels above IP. Just as routers interconnect arbitrary media, the net can be made of gateways that interconnect arbitrary media characteristics and application goal regimes. The breadth of media transitions are getting wider - cell to dense wireless up to all-optical terabits. The breadth of application goals with convergence is also getting wider - individualized media to mass broadcast, damage-tolerant to six nines reliability. An architecture that aims to accommodate and make appropriate transitions for the ranges that might be expected in the net of 2005 or 2010 could be called a diversinet.

1 Introduction

The recent growth of the Internet has been phenomenal in many ways, and shows the broad vision of its creators. However, when we consider that growth in more details, we see significant changes in both the network infrastructure and in the applications supported by the network. These changes are stressing the Internet model beyond its original design.

In this paper we attempt to elucidate some of those changes, and consider the problems posed by the increasing heterogeneity of network services and applications, in particular the convergence of data, voice and televisual networks and applications.

One possible approach by which this heterogeneity can be tamed and taken advantage of, is to embrace diversity at the levels above IP. Just as routers interconnect arbitrary media, the network can be made of gateways and middle-boxes that interconnect arbitrary media characteristics and application goal regimes. The breadth of media transitions are getting wider – cell to dense wireless up to all-optical terabits. The breadth of application goals with convergence is also getting wider – individualized media to mass broadcast, damage-tolerant to six nines reliability. An architecture that aims to accommodate and make appropriate transitions for the ranges that might be expected in the net of 2005 or 2010 could be called a diversinet.

2 The curse of heterogeneity

The traditional world of the Internet, supporting email, FTP, web and little else is coming to an end. We are already seeing the convergence of the PSTN with the Internet (e.g. 3G cellular systems are expected to use IPv6 as their transport; many Telcos are rolling out VoIP backbones) and the future holds the imminent convergence of the broadcast TV and cinema industries with the Internet.

This convergence blurs the distinctions between different classes of application and network. The traditional, separate, networks have provided radically different quality of service, transport (packet vs circuit, for example) and reliability for these different applications. Accommodating them in a single network will prove to be a challenge.

In the remainder of this section, we discuss the network of the future: in terms of applications it will support, the infrastructure that will support them, and their addressing/naming requirements.

2.1 The applications

Envisaging the applications which may run over the next generation network, we see a broad spectrum of requirements.

- Traditional batch oriented applications – such as email, netnews – provide relatively short flows, with elastic delivery bounds, and little need for enhanced QoS. They work well on best-effort networks and care little for high bandwidth or low latency.
- Traditional interactive applications – such as the world-wide web – also provide relatively short flows, but with significantly tighter latency bounds. Bandwidth demands are not typically high, although many users use high bandwidth as an approach to achieving low latency.
- Traditional bulk transfer – such as FTP, and larger web downloads – need a high average bandwidth, but are not unduly concerned with latency or with short-term fluctuations in bandwidth.
- Streaming audio/video applications – such as RealAudio, Windows Media, TV distribution via IP – need consistent bandwidth, but care little about delay (in most cases). They are also tolerant to loss, provided that loss is predictable.
- Interactive audio/video applications – voice over IP, video-conferencing, – and interactive network games need low latency and, in some cases, relatively high bandwidth. They are also relatively insensitive to loss, since their latency bounds do not allow time for recovery.
- Interactive non-audio/video applications – such as instant messaging – are sensitive to latency, but low bandwidth.
- Sensors and other devices with extremely limited resources. These are very tolerant to damage, but need to be trivial to deploy and need to scale massively and self-organize. Fate sharing between devices must be reduced, which implies a stateless transport.
- ‘Bluetooth’ devices, such as web-pads, appliances, etc. Need to be simple, autoconfiguring, mass produced. They can have a range of requirements: for example wireless headphones need low jitter and loss, but a wireless TV remote can easily cope with this.

It is interesting to note that we are beginning to see applications which are tolerant to loss, rather than requiring perfect transport. Most of the developing audio/visual applications can tolerate packet loss, still others can use corrupted packets, a feature which many data networks do not support.

Furthermore, many emerging applications are strongly multiparty, imposing the additional requirement of efficient multicast support, and leading to additional heterogeneity.

2.2 The network

In parallel to this application development, the network has also undergone significant and rapid evolution: optical networking has led to a massive increase in backbone capacity, DSL and cable modem technologies have led to significant improvements in last-mile access rates, and wireless networking has become a reality, both wide-area (e.g. GSM) and local-area (IEEE 802.11). These trends can be expected to continue, and accelerate over the next few years.

With the introduction of DWDM technology, the bandwidth available in optical networks is expected to increase still further, widening the gap between the line-speed of the network and the speed of the (electronic) switching control plane. Such networks will deliver high bandwidth with very low latency and loss rates, but dynamic setup and control of these networks is difficult, and will likely lead to circuit switched solutions.

The success of optical networking may well lead to the displacement of wired access in all but the local area. In five years, it is likely that many SMEs will have optical access (many large businesses have it now...) and that DSL and cable modems will be beginning to be supplanted by fibre to the home.

What is unclear is the effect this will have on the backbone. It is certainly feasible that providers will be able to install enough fibre, and light enough wavelengths, that the backbone will remain uncongested. However, it is also possible that the roll-out of optical technology to the edges of the network will lead to the return of congestion in the backbone. Something to beware of, especially if the backbone is effectively circuit switched by that time.

The other significant trend in the evolution of the network is wireless access. Already, cellular technology is widespread and the first generation PDA/phone hybrid devices are appearing on the market. These are expected to proliferate, leading to a shift in the dynamic of the network: these devices will be 'always on' but be mobile, with widely varying access characteristics depending on radio reception. The convergence of wireless voice and wireless data will also push the need for differentiated service in the wireless network.

We expect that low-bandwidth wireless access will be ubiquitous in five years, many devices having permanent connection where-ever they roam. In addition, high bandwidth local area wireless will be commonplace (either 802.11 or Bluetooth based, or more likely a higher bandwidth successor to these).

It is important to note that these different technologies have widely varying characteristics: optical is very high bandwidth, low delay, virtually zero loss; wireless is significantly lower bandwidth, higher delay and lossy (both packet and bit errors). Optical networks are well suited to circuit switched control and long lived flows, wireless works well with packet switching, ephemeral traffic flows, and mobility.

2.3 The address space

One of the fundamental features of the Internet Protocol suite is that it provides a uniform network interface and address space. This allows for an application running on one machine to communicate with any other machine on the network, with the same syntax no matter where the other machine is.

Unfortunately, this is changing rapidly. The major driving force for this change is the lack of IPv4 address space, meaning that it is not possible for many organizations to obtain globally unique network addresses. This has resulted in significant use of private addresses, and the deployment of network address translation (NAT)¹. This has resulted in a fragmented and brittle network: NAT can require detailed knowledge of the contents of the packets, especially for protocols such as SIP which are used to setup other traffic flows, and which can have significant security implications.

The lack of address space has also led to the development of alternatives such as IPv6, which extends the 32 bit address space to 128 bits. This is, in theory, sufficient for a return to a single global address space. With the deployment of 3G wireless telephony systems, based on IP, it is likely that IPv6 will see significant deployment over the next few years.

The transition to IPv6 will be painful, however. There are millions of IPv4 devices, and communication between the new systems and the old will have to be facilitated: this implies some form of NAT at the boundary.

We therefore see the use of NAT increasing over the next five years, to the extent that many – perhaps most – communications will pass through some form of NAT (either between IPV4 realms or across the 6-to-4 boundary).

3 Implications

The trends we have identified have a number of implications for future networks and for future networked applications.

As noted previously, we expect the use of NAT to increase over the next five years, and that middle boxes (NAT, firewalls) will be common-place. This has the implication that pure end-to-end IP connectivity will vanish for most applications, and that almost all traffic will pass through some form of middle-box.

This, in turn, has implications for security (end-to-end authentication is difficult when NATs must modify packets in flight), addressing (since there is no uniform IP address space) and transport (protocols should be designed to avoid use of embedded addresses, to make translation easier).

The increasing range of network technologies leads to further divergence: optical and wireless networks have very different characteristics, and the ideal transport protocol for one class of network is badly suited to the other. The implication is that we will see protocols optimized to the low level transport, and indeed we already see divergence in the protocols used in different parts of the Internet (compare the use of MPLS flow switching for high bandwidth optical networks, with loss tolerant header compression and error tolerant transport in wireless networks) and this trend can only continue. In the extreme, protocol translation may be required to transit the boundary between the different classes of network.

Scalability of protocols and systems is becoming more of an issue, both in the range of devices, and the number of end-points, which must be supported. The implication is that content providers will produce different content for the different classes of device. Two examples illustrate this:

- A number of website are produced in multiple editions, depending on the class of device on which they will be presented. Partly this is to work around bugs and limitations of certain browsers, but it can also be to work around limitations of the presentation media: sites which have two versions, one optimized for viewing on a PC with high bandwidth connection, the other suited for a PDA with low bandwidth wireless access, are becoming common.

¹NAT tries to preserve the appearance of a single uniform address space by rewriting the addresses of packets in transit.

- Providers of streaming audio/video services routinely produce multiple versions of their content, to suit different speed connections and the limitations of different end-systems.

This problem becomes more acute as the number of receivers increases: not only does a server need to generate multiple versions, but unless robust multicast support is present, it must generate multiple copies of each version. This is the “N-squared” problem: grappling with content transcoding/transformation so that providers can manage the scaling, even with heterogeneity.

Multicast, once again, leads to more requirements being placed on middle boxes. In particular, with the increasing heterogeneity of the network, the multicast ‘routers’ may need to perform rate adaptation and congestion control, error correction, authentication, etc. The initial stages of this are being observed with the current deployment of content distribution networks.

Finally, there is the implication on feedback and control. In the current Internet, TCP/IP has an implicit feedback loop which treats loss as an implicit congestion signal, and real-time transport has a background reception quality reporting protocol with similar assumptions. Neither of these are valid once links with non-congestive loss (e.g. wireless) are present in the network. Further, there are no widely agreed upon means of providing explicit congestion control or loss feedback for large scale multicast flows – important for the future scaling of the network.

4 Key issues and challenges

The key issue for future network research will be *overcoming heterogeneity*, to provide a *semantically consistent* network interface.

What do we mean by this? The traditional approach has been to develop a single protocol layer which abstracts the lower layers, providing an internetwork function – “one protocol to bring them all, and in the darkness bind them” – and end to end connectivity.

There are a number of benefits to the traditional model: simplicity, consistent interface, uniform service, and single address space. There are also a number of significant limitations, which we have noted: poor handling of heterogeneity, and limited support for advanced services (e.g. quality of service, security, multicast). Why then has IP been such a success? Because it exposed the underlying network transport of the time.

The next generation approach is to preserve that model at a semantic level, but to embrace the presence of translation devices and the lack of pure end to end transport, to provide a consistent interface to the disparate network transport services, whilst exposing details of the transport path to those applications which need this information. Recognize that the network of the future will be more than a single, best effort, packet delivery system, and future proof the design by embracing heterogeneity when necessary.

What do we need to focus on to achieve this?

Naming In a world where devices do not necessarily have a unique address, or where they may have an address which is not necessarily directly reachable from a particular point in the network, it will be necessary to extend the directory services available. The DNS works well for a single address space, but will not support multiple, semi-overlapping, address spaces without significant overhaul.

Protocols Do packets make sense in a world where we have ultra-high rate optical flows which are hard to switch? Do circuit switched flows make sense in the wireless world, where connections may be lost at any time (due to mobility, fading, interference, etc)? Do we need a hybrid protocol?

Protocols Is the end-to-end IP model still viable? If so, how do we return to a world where it is feasible? How much of the functionality provided by middle boxes might we wish to keep?

Scalability How can we leverage the presence of middle boxes to enhance the scalability of our systems? Multicast routing is the first stage, but also in terms of congestion control, transcoding, error correction, security?

Reliability IP has traditionally provided best effort delivery of packets, with delivery failure resulting in the complete loss of the packet. What are the implications of partial delivery: delivery of packets with corruption? How does this affect security?

Security How does the presence of middle boxes affect security? The necessity of modifying data en-route invalidates many of the traditional approaches, and new techniques will be needed.

The challenge, therefore, is to build a heterogeneous network which can work as an integrated entity. It is no longer sufficient to hide the strengths and weaknesses of the underlying networks – rather they must be exposed, allowing smart applications to optimize their delivery of data to the needs of the network. Application level framing taken to the extreme.

5 Research directions

How do we achieve the network of 2010 that embraces diversity and new structures? The enablers include:

Highly shared node resources for architecturally focused network and distributed systems researchers (a concept that one of the authors collaborating with Tom Anderson of University of Washington has proposed as the ACCESS infrastructure). This resources should be node clusters in multiple regions of widely diverse media (optical, wired, fixed and mobile wireless of heterogeneous technologies), of which commercial Internet media are a part - this testbed should be far more virtual than existing testbed, and ride on the natural curve of diversification that is occurring operationally.

Further, the transition from IPv4 to IPv6 has largely been viewed as an engineering exercise, devoid of research problems. In practise though, it is just the first of a range of interworking problems which needs to be solved. By actively driving this transition, and by actively driving the deployment of mechanisms for QoS, “multicast”, mobility, security, etc., it will be possible to drive the development of the next generation of networked systems.