

Internet Big Bang Theory

Curtis Villamizar <curtis@avici.com>

April 1, 2000, updated November 13, 2000

Abstract

Cosmologists have now accepted the notion of a big bang theory as the beginning of the universe. A leading contender among theories is now the inflation theory in which the universe briefly expands at speeds orders of magnitude greater than the speed of light. The Internet and its appetite for bandwidth is currently expanding at a rate that until recently economists considered impossible. The notion that the Internet would overtake the voice network, first in terms of bandwidth usage, and then in terms of profit and revenue, is now well accepted. Some of the theories about the Internet's expansion in the coming 2–5 years seem analogous to cosmology's inflation theory in that prior notions of the laws of physics must be suspended or modified in order to support the theories.

Despite initial impressions that one may draw from the attempt at humor in the way of an unusual analogy, this paper provides a serious look at the problem of scaling the Internet by one or more decimal orders of magnitude. Protocol scalability is examined, considering potential roles of optical switching and the new class of terabit routers.

Contents

1	The Internet Universe So Far
2	Expansion of the Universe
2.1	The All Optical Backbone
2.2	Making Use of Both Packet Routing and Optical Switching
2.3	How Big Can We Make This
3	Conclusions

Introduction

The analogy between cosmology and engineering Internet Service Provider (ISP) backbones may be stretched a bit in this paper. Possibly more than just

a bit. The similarity is that expansion is rapid, even unbelievably so. The inspiration for this analogy was a keynote speech delivered by a venture capital firm where an expansion in backbone capacity of 10^5 to 10^6 over the next 5 years was predicted.

Cosmologists peer into the past and gather evidence as to an amazing transformation in the form of an enormously rapid expansion of the universe that appears to have occurred billions of years ago. Venture capitalists and the financial community peer into the future at what they expect to be an enormously rapid expansion of the Internet, and their fortunes. At times the financial community seems intoxicated with the scale of their own predictions. Engineers take a far more sober view and attempt to peer into the future, keeping in mind the “laws of physics” that pertain to their universe and trying to figure out how they are going to make it all work.

1 The Internet Universe So Far

The Internet has already expanded by about 10^6 in under 15 years and has certainly presented challenges along the way. In 1986 the Internet in the US was dominated by a few 56 Kb/s backbones operated by US government agencies, most notably the NSF. At that time the Internet barely existed outside the US. By 1988 the NSFNET had a 1.5 Mb/s backbone and by 1992 the NSFNET was upgraded to 45 Mb/s. The Internet was already growing in the US and worldwide. By 1995 there were multiple 45 Mb/s backbones. Among these were ANS, InternetMCI, Sprint-Link, and UUNET. By year 2000 there are even more Internet backbones, the largest among them using multiple 2.5Gb/s circuits between pairs of routers and numerous plans to upgrade to 10Gb/s circuits. The range in size of any single backbone is 5×10^4 to about 10×10^9 over the period of 1986 to 2000. Over the past 5 years growth was from about 5×10^7 to about 5×10^9 or 10^2 in 5 years.

During the period from 1986 to 1994 the largest Internet backbone was the US NSF funded NSFNET, an enormously successful research project. The

routers were cobbled together by today's standards. These routers used off the shelf engineering workstations with increasingly specialized packet forwarding cards. Commercial routers that put an increasing amount of packet forwarding support in hardware appeared in the 1994 to 1995 time frame.

Many of the lessons learned in the NSFNET research project were apparently forgotten or ignored in the initial "high end" commercial routers. The result was a very fragile Internet circa 1996 to 1997, one that was barely able to accommodate its own growth.

From 1995 to 1998 router performance improved only incrementally. In order to use higher speed interfaces and in order to accomplish traffic engineering goals, ISPs reluctantly turned toward ATM backbones. The scaling properties of IP over ATM were less than ideal and well known [6]. There were initial objections to ATM encapsulation efficiency, the "cell tax" argument, but ATM Segmentation and Reassembly (SAR) and the purely connection oriented nature of ATM proved to be the key limiting factors.

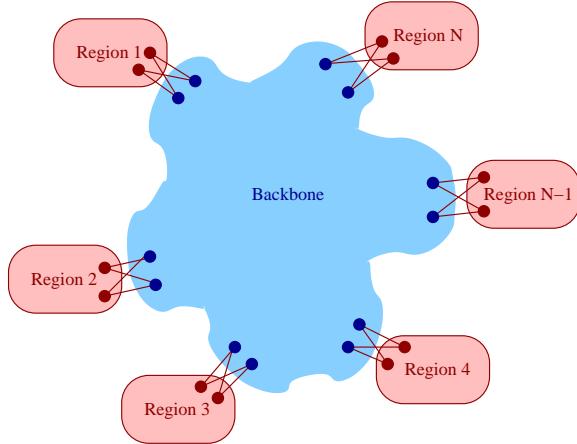


Figure 1: The Cosmic Background: Some of the better engineered Internet backbones are subdivided into regions in order to improve scalability properties.

In 1997 and 1998 a new generation of the Internet routers appeared. These routers were capable of supporting first OC-12c (655 Mb/s) and later OC-48c (2.5 Gb/s) interfaces. The first interfaces were not capable of forwarding at anywhere near full line rate, particularly if the routers were fully populated with high speed interfaces. Later these routers supported multiple OC-48c interfaces at full line rate. This generation of equipment is known as "gigabit routers" for their aggregate capacity in the range of tens to hundreds of gigabits per second.

In addition to the problem of making routers go faster was the problem of putting them together in

some sensible fashion. In the better designed Internet backbones the topology was subdivided into regions as shown in Figure 1. There are typically on the order of ten regions. The backbone has only tens of routers and/or switches. In the larger backbones, these nodes (router or switch) are interconnected via OC-48c with occasional multiple parallel OC-48c between nodes. Within a region there are on the order of 100 routers. This design avoids a set of problems known as collectively as "interior gateway protocol (IGP) scaling issues".

The links state protocols OSPF and IS-IS are used by all Internet service providers of any size. These are known as Interior Gateway Protocols (IGP). The SPF calculation used in link state protocols such as OSPF and IS-IS required order(N^2) computation or with an optimization applicable to less densely meshed networks order($N \log_2 L$) where N is the total number of nodes and L is the total number of links in a IGP area. With even 100 nodes, a full mesh network required on the order 10,000 operations or 1,000,000 operations for the SPF calculation.

If a network is fully interconnected, then reflooding problems can become quite severe. When a link goes down, routers on each side advertise the change to each of about N adjacencies. They in turn advertise to each of their N adjacent peers, not knowing whether the peer already knows of the change. Each router receives up to N copies of each advertisement. If a router goes down, each adjacent router announces a loss of adjacency to its $N-1$ remaining peers. Each peer advertises this again $N-2$ times. Each router gets $N-1$ copies of $N-1$ announcements. With 101 nodes, each router can receive some 5,000 to 10,000 route announcements of which 100 are unique, resulting from a single event.

For each advertisement the router must decide whether to reflood. In a full mesh this involves searching the N^2 set of advertisements. A naive IGP implementation might use an $order(N)$ search of the link state database to find an advertisement and see if the latest information is already known to determine whether to reflood. With 101 nodes, 10,000 advertisements must be processed and the naive search would be expected to yield 50,000,000 comparisons. An $order(\log_2 N)$ search yields 160,000 comparisons.

As IGP sizes grow to hundreds of routers, use of a full mesh becomes a severe problem. Figure 2 illustrates a variation of a full mesh, a dual plane design. The dual plane design allows two backbone routers per region, eliminating a single point of failure, without doubling the full mesh. This design presents no problem at all for the modest backbone sizes of this type in use today, which typically have a few tens of

routers in the backbone.

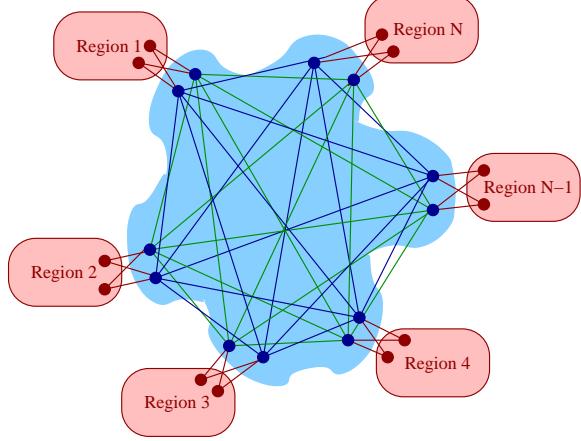


Figure 2: The Full Mesh Backbone: Within a backbone which is designed using IP over a circuit switch technology such as ATM, full mesh connectivity or dual plane full mesh is used. For N regions each router has about N direct logical adjacencies. Substantially increasing N is very difficult due to scaling properties of link state routing protocols.

Along with the later “gigabit routers” has come further improvement in protocol design. The use of the Multi-protocol Label Switching (MPLS) [4, 5, 8], POS (PPP or Packet over SONET) [10], and “Constraint Based Routing” [1] support comparable or better traffic engineering capability than what provided by ATM and eliminate the “IGP scaling” problems.

In an IP over MPLS over POS design the number of logical adjacencies per router is equal to the number of physical adjacencies, which is relatively constant. With no improvement in the raw performance of the routing protocol engines the elimination of the full mesh reduces the computation to order($N \log_2 N$) and reduces the readvertisement problems to an order(1) issue. This design is illustrated in Figure 3. With all other factors equal, meaning no improvement in the raw performance of the routing protocol engines, backbones IGP areas about 5 times larger should be possible. The number of regions can therefore be increased if the physical topology is conducive to such a change (it may not be).

Within the regions themselves there are typically under 100 routers. It would be safe to put this figure at on the order of a hundred for the largest regions of the largest Internet providers. Two or more routers within a region are adjacent to the backbone. These are known as “interconnect routers” or alternately “hub routers”. These interconnect routers

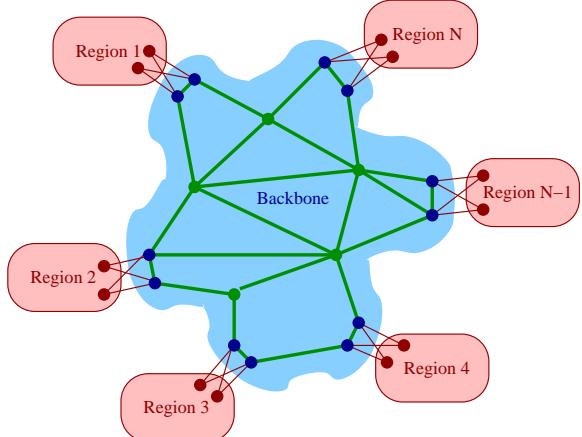


Figure 3: Eliminating the Backbone Full Mesh: The MPLS protocol is now widely accepted as the preferred means to keep the number of adjacencies relatively constant while allowing the number of regions to increase substantially.

are logically adjacent to each of the edge routers known as “access” or “aggregation” routers. The mapping of physical connectivity to logical connectivity within the region is generally based on either ATM or SONET Time Division Multiplexing (TDM).

Figure 4 illustrates such a design. The logical topology with two routers adjacent to the backbone is described as a “dual star” topology. Some inefficiency is introduced by requiring intra-region traffic to always cross one of the few interconnect routers within the region.

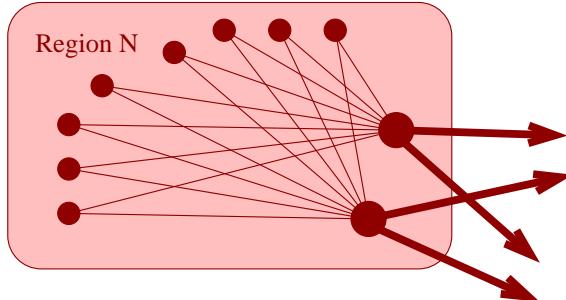


Figure 4: Full Mesh Regions: Analogous to the full mesh backbone problem in Figure 2 is the problem of supporting a high number of adjacencies within a region designed using switch technology or SONET. Scaling is improved by building a logical dual star topology rather than a full mesh but is still limited.

Scaling within a region can also be improved using MPLS. In Figure 5 a region is shown with MPLS routers replacing the ATM switches or SONET Add-

Drop Multiplexors (ADMs). The region is often physically interconnected as one or more regional fiber rings with second tier cities attached to the rings at at least two places on the ring.

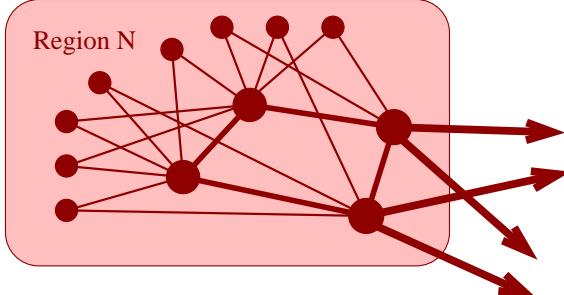


Figure 5: Eliminating the Region Full Mesh: MPLS can also be used within a region to improve scalability, allowing the region to include a larger number of second tier cities.

2 Expansion of the Universe

As the Internet expands by decimal orders of magnitude every few years, new generations of equipment and significant changes in network design are required. The process has been described as just completing deployment of one generation of network in time to begin planning its successor. Many ISPs plan one or two generations beyond the current network in order to insure that appropriate equipment exists and the network build logistics are anticipated.

Detailed arguments supporting a continued demand for bandwidth are beyond the scope of this paper. Some of the (more credible) arguments are stated briefly below.

1. Consumer access speeds are increasing by approximately 1-2 orders of magnitude from 30-50 Kb/s to 0.5-2 Mb/s (DSL and cable modem). The number of consumers using the Internet and the amount of use of the Internet per user is increasing. As bottlenecks in connectivity from consumer oriented Internet providers to backbone providers are removed, this trend will be seen as growth in backbone traffic.
2. The access speeds at which enterprises connect to the Internet is increasing even more dramatically than the the increase in consumer access speeds. Medium to small businesses are connecting in larger numbers. Use of the Internet is transforming from static web pages that serve

as online posters or brief brochures to detailed product information and online commerce.

3. The use of graphics has gone from 5-10 KB icons to low resolution images. As bandwidth becomes available the use of higher resolution images becomes more common. Other media such as audio and animation become viable as bandwidth becomes available. Applications such as two-way voice and video-conferencing are certainly in demand and may soon become viable.
4. Inventive minds are certain to come up with applications that have yet to be considered.

Venture capitalists, economists, engineers and others will certainly argue over the number of orders of magnitude of growth that lie ahead, the rate of growth, and when or if growth will level off. It seems fairly certain that in the immediate future at least 1-2 orders of magnitude of further growth will occur.

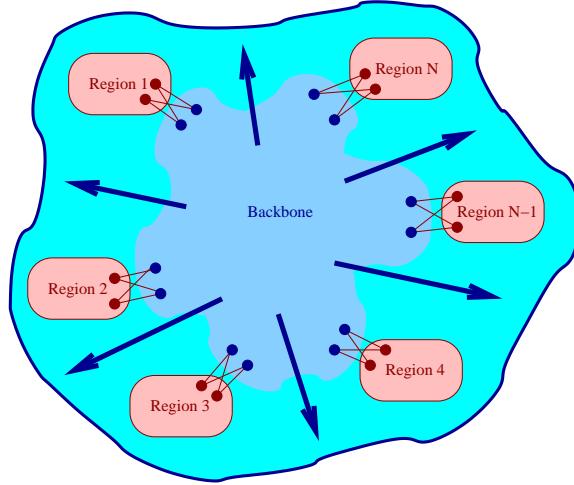


Figure 6: The Big Bang: Predictions of decimal orders of magnitude growth in the Internet will need to be accommodated with a larger backbone and therefore more regions and larger regions. The speed of interconnections within the backbone will certainly challenge the limits of technology. Some questions have been raised regarding the roles of optical and electronic switching in future backbones.

To an engineer what matters is how to accommodate this growth. A key issue is where in the design to accommodate the growth. Figure 6 draws on our initial analogy and suggests that a common theme is expansion of the backbone to encompass what is now the regions and the use of denser regions outside this backbone. No public designs have come forth but some possible hints have been provided. The notion

of “all optical” backbones has received significant attention. In the next section the credibility of this idea is questioned.

Though all optical backbones may or may not prove feasible, optical network elements will play an important role in future backbones. MPLS is being proposed as a control protocol for optical backbones. There is a great deal of activity in this arena [7, 11, 13, 3, 12, 2].

2.1 The All Optical Backbone

The notion of an all optical backbone consisting of an optical core surrounded by gigabit routers seems seriously flawed unless some undiscovered or not yet perfected technology allows a much finer multiplexing granularity. Figure 7 shows an expanded backbone which a region attached to it. The large pipe into the backbone with smaller tributaries is an attempt to illustrate a connection to the optical backbone which is broken down into smaller flows to a large number of other regions.

Assume for the moment that IGP scaling issues are somehow avoided and optical paths signaling is accomplished somehow (possibly using out-of-band MPLS control plane). Even at the current scale with on the order of 1,000 access routers, each access router would require an adjacency to the 1,000 other routers and therefore would have to support enough interfaces to terminate 1,000 wavelengths, presumably at OC-48c to OC-192c speed. This is certainly well beyond the scale of a gigabit router.

The only potential candidate to save this design would be some technology that would allow multiplexing within a wavelength. It has been suggested that optical Time Division Multiplexing (TDM) may be feasible if switching times were faster. To achieve today’s size of 1,000 access routers, 100 time slots on 10 wavelengths could be used. Since delays of over 100 msec are considered somewhat excessive, time slots would have to be spaced less than 1 msec apart. The best optical switching times are currently on the order of 10 msec.

TDM has been criticized as highly wasteful of bandwidth. This approach, if it were to work in the future, would still only make sense if the savings realized by using optical switching outweighed the cost of requiring additional wavelengths to make up for wasted bandwidth.

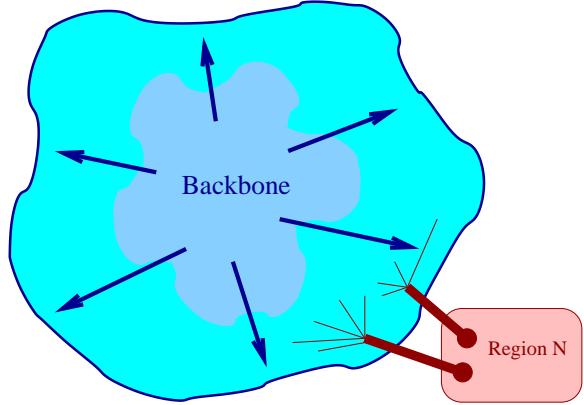


Figure 7: The All Optical Backbone: The all optical backbone is problematic from the standpoint that the region routers on the edge of such a backbone must maintain direct logical adjacencies with at least one router every other region. If there are N other regions and bandwidth is allocated in OC-48c or OC-192c wavelengths, then each region router must have N such interfaces. The only all optical alternative is optical TDM but at optical switching rates orders of magnitude higher than achievable today.

2.2 Making Use of Both Packet Routing and Optical Switching

Optical switching will certainly have a place in future designs, but all optical backbones do not seem feasible. Electronic multiplexing of data will still be required closer to the core of the network. The best candidate for this remains IP routers, only of a much larger scale than today’s gigabit router.

In the previous section an extreme position was postulated, that of an all optical backbone. Some design implications were examined to see if any technology on the horizon could accomplish this. This is not a typical engineer approach, but the extreme position has been postulated in public and is gaining some attention so it was worth examining.

A better approach is to examine what the introduction of emerging technologies can do to improve on the existing working design, evolving the backbones toward larger scale. For example, MPLS can allow the backbone size to increase relative to an overlay model such as that used in IP over ATM backbones. Another well accepted technique is the use of IGP areas. Through the use of areas an IGP may be able to grow to something on the order of 10,000 nodes or more rather than 1,000.

The logical backbone may expand to a larger topology which encompasses what is now the major regional fiber rings and the regions in the logical topol-

ogy. Beyond today's access routers, the Internet is a star topology, with circuits back-hauled from the Point of Presence (POP) where the access routers are located directly to customers (including medium sized or smaller dial and DSL access providers who are customers of the major ISPs). In the future there may be a much larger number of geographically smaller regions encompassing metropolitan fiber rings and including POPs in second tier cities and in more remote places. More aggregation will occur in these outlying places as Internet penetration continues.

Figure 8 suggests a tiered approach in which there remains a relatively small core. This may be optically switched with terabit routers around the perimeter of such a core. This core may be the backbone area of the ISP's IGP. The term "backbone area" is an OSPF term meaning central area and is always numbered area 0. The non-backbone areas (using OSPF terminology, meaning areas other than area 0) may also be optically switched with terabit or high end gigabit routers closer to the edge.

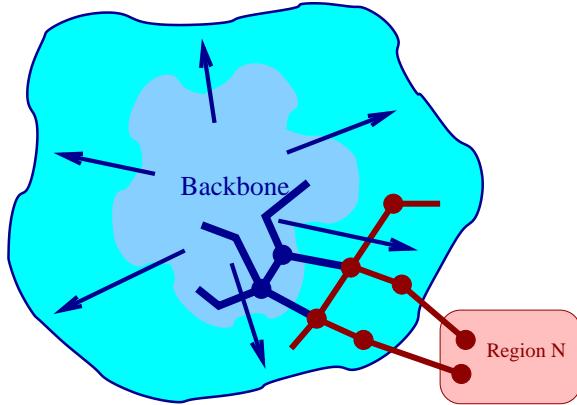


Figure 8: The Mixed Optical and Electronic Backbone: Dense Wave Division Multiplexing (DWDM) capable of supporting hundreds of wavelengths and terabit routers capable of supporting hundreds of interfaces and concatenating large numbers of physical interfaces into smaller numbers of logical interfaces exist today. Combining these in a mixed optical switched and electronic routed core seems the most likely candidate to accomplish further Internet scaling. The backbone may have to be subdivided into areas, using area support provided in the existing link state protocols.

The terabit routers in the core are required to achieve multiplexing that is accomplished via packet routing. This allows the area 0 routers to multiplex traffic from the 1,000 to 10,000 routers closer to the edge, allowing those routers to terminate only a few

wavelengths. The area 0 routers would be required to terminate a very large number of wavelengths.

2.3 How Big Can We Make This

It is useful to examine both the scaling limits of a mixed routing and optically switched network and how such a network might evolve as each decimal order of magnitude of growth occurs.

First we must consider the current size of Internet backbones. Global Crossings has been very open about their design [9]. It utilizes tens of backbone routers interconnected primarily with OC-48c with some OC-12c. The regions are interconnected primarily with OC-12c. UUNET and Sprint have been somewhat less open about their topology, reluctant to provide maps but hinting that nodes are interconnected with multiple OC-48c interfaces. Four or more OC-48c interfaces are used in places, providing the equivalent of OC-192c rate or greater. UUNET has recently announced intention to build an OC-192c based backbone using routers capable of supporting 16 OC-192c interfaces.

If the existing design were scaled by one decimal order of magnitude, the largest backbone routers supporting on the order of a hundred OC-48c interfaces would be required, or two dozen OC-192c interfaces. Access routers supporting a few OC-48c interfaces would be needed. If ATM or SONET gear in the access networks were replaced with routers, these might require a few dozen OC-48c interfaces or 5 or 10 OC-192c interfaces.

Optical switching may play a role in both the backbone and within the regions, with terabit routers providing the multiplexing within the backbone that cannot be provided by optical technology. Optical switching is unlikely to replace the backbone because to do so would require a very large number of lambdas at the routers at the periphery and aggravate IGP scaling problems.

If the network were scaled a second order of magnitude, a small number of core routers would require on the order of 1,000 OC-48c or on the order of 250 OC-192c. Optical switching within such a core would reduce the number of these routers or the number of interfaces per router but could not entirely replace them. At this point IGP areas are almost certain to be required as are MPLS techniques to establish paths across areas. The access routers would require on the order of a hundred OC-48c interfaces or tens of OC-192c interfaces with smaller terabit routers displacing the gigabit routers. The backbone would almost certainly have to be pushed out to the boundary of the current major regional fiber rings with more

regions to keep the number of nodes with regions sufficiently bounded.

A third order of magnitude would require the use of OC-768 in the backbones, on the order of 500-1,000 such interfaces if optical switching were not used. Scaling to this size may require pushing the IGP backbone area farther out to allow optical switching only within the very inner core. Alternately a larger number of the core terabit routers could be used but this may be less cost effective. The access routers at this point would also have to be multiple bay terabit routers.

Keep in mind that the scaling of the single provider largest providers has been considered. The major providers would all be expected to be scaling their network in size. The technology to interconnect providers of this magnitude would have to include the very high end terabit routers for the same reasons of flexibility in terms of multiplexing.

Currently available terabit routers have been designed for massive scaling of the type described here but routers of this size (on the order of a thousand interfaces) have yet to be demonstrated. It is likely that these terabit routers will be deployed early in the process, during the first order of magnitude of scaling, in single bay configurations as backbone routers. As growth continues optical switching equipment will be introduced in addition to upgrades to the backbone terabit routers. Terabit routers will displace gigabit routers closer to network edges as those closer to the core are upgraded to multiple bay configurations and/or upgraded to OC-192c. Going from the second to third order of magnitude of scaling requires the maximum configurations of terabit routers, and the use of multiple fibers, Dense Wave Division Multiplexing (DWDM) and optical switches between entities.

3 Conclusions

The technology brought to bear on this problem is certain to include terabit class routers, essentially more of the same technology, only much bigger. Such routers exist today in designs scalable to hundreds of OC-48c interfaces and within a year or two on the order of 1,000 OC-192c through card swaps in the existing chassis. Such chassis have been designed to support up to OC-768 interfaces.

Optical DWDM technology is clearly required and is in fact a key enabler to the continued growth of the Internet. Optical switching also provides clear benefits in terms of switching at a reduced cost but is not likely to be the panacea that some have implied

or claimed.

The best technology currently on the horizon would very likely be pressed to its limits to accommodate three orders of magnitude of growth. This technology would be unable to accommodate growth beyond three orders of magnitude. Currently announced products would certainly not go even this far. Even going to three orders of magnitude requires the largest terabit routers currently being designed, multiple fibers each filled to capacity with DWDM equipment capable of densities on the upper limit of technology reported to be demonstrated in research environments and a great deal of optical switching equipment running at the limits of announced research in that field.

Projections of five to six orders of magnitude of growth in Internet capacity or data networking capacity in five years motivated the analogy to cosmology used in this paper's introduction. Perhaps the analogy to theories of a universe expanding orders of magnitude faster than the speed of light is a good one. It is also entirely possible that both will prove accurate.

While there are many uncertainties, two things seem very certain in the near future. Cosmologists, particle physicists and others will be very busy trying to interpret the clues that nature has provided about the origins of the universe. Optical engineers, ASIC designers, software engineers, and network designers will be very busy trying to produce designs and equipment that stay ahead of the growth in demand for the Internet and other uses of network bandwidth.

References

- [1] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus. Requirements for Traffic Engineering Over MPLS. Technical Report RFC 2702, Internet Engineering Task Force, 1999. <ftp://ftp.isi.edu/int-notes/rfc2702.txt>.
- [2] Lou Berger, Y Rekhter, and Kireeti Kompella. Link Bundling in MPLS Traffic Engineering. Internet Draft (Work in Progress) [draft-kompella-mpls-bundle-03](ftp://ftp.isi.edu/internet-drafts/draft-kompella-mpls-bundle-03.txt), Internet Engineering Task Force, 9 2000. <ftp://ftp.isi.edu/internet-drafts/draft-kompella-mpls-bundle-03.txt>.
- [3] Lou Berger, Hal Sandick, Y Rekhter, John Drake, Debanjan Saha, Kireeti Kompella, Debasish Basak, Jonathan Lang, and Krishna Mitra. Link Management Protocol (LMP). Internet Draft (Work in Progress) [draft-ietf-mpls-lmp-01](ftp://ftp.ietf.org/internet-drafts/draft-ietf-mpls-lmp-01.txt).

lmp-00, Internet Engineering Task Force, 9 2000. <ftp://ftp.isi.edu/internet-drafts/draft-ietf-mpls-lmp-00.txt>.

[4] Ross Callon, George Swallow, N. Feldman, A Viswanathan, P. Doolan, and A. Fredette. A Framework for MPLS. Internet Draft (Work in Progress) draft-ietf-mpls-framework-05, Internet Engineering Task Force, 9 1999. <ftp://ftp.isi.edu/internet-drafts/draft-ietf-mpls-framework-05.txt>.

[5] Ross Callon, A Viswanathan, and E. Rosen. Multiprotocol Label Switching Architecture. Internet Draft (Work in Progress) draft-ietf-mpls-arch-07, Internet Engineering Task Force, 8 2000. <ftp://ftp.isi.edu/internet-drafts/draft-ietf-mpls-arch-07.txt>.

[6] R. Cole, D. Shur, and C. Villamizar. IP over ATM: A Framework Document. Technical Report RFC 1932, Internet Engineering Task Force, 1996. <ftp://ftp.isi.edu/in-notes/rfc1932.txt>.

[7] Rob Coltun, Y Rekhter, John Drake, and Daniel Awduche. Multi-Protocol Lambda Switching: Combining MPLS Traffic Engineering Control With Optical Crossconnects. Internet Draft (Work in Progress) draft-awduche-mpls-te-optical-02, Internet Engineering Task Force, 7 2000. <ftp://ftp.isi.edu/internet-drafts/draft-awduche-mpls-te-optical-02.txt>.

[8] Dino Farinacci, Tony Li, A. Conta, Y Rekhter, Dan Tappan, E. Rosen, and G. Fedorkow. MPLS Label Stack Encoding. Internet Draft (Work in Progress) draft-ietf-mpls-label-encaps-08, Internet Engineering Task Force, 8 2000. <ftp://ftp.isi.edu/internet-drafts/draft-ietf-mpls-label-encaps-08.txt>.

[9] Alan Hannan. Deploying a greenfield network using pos and mpls. In Stan Barber, editor, *North American Network Operators' Group Meeting Notes*, San Diego, California, May 1999. North American Network Operators' Group, NANOG. <http://www.nanog.org/mtg-9905/hanna.html>.

[10] A. Malis and W. Simpson. PPP over SONET/SDH. Technical Report RFC 2615, Internet Engineering Task Force, 1999. <ftp://ftp.isi.edu/in-notes/rfc2615.txt>.

[11] Y Rekhter, John Drake, Daniel Awduche, and Debasish Basak. Multi-protocol Lambda Switching: Issues in Combining MPLS Traffic Engineering Control With Optical Cross-connects. Internet Draft (Work in Progress) draft-basak-mpls-oxc-issues-01, Internet Engineering Task Force, 2 2000. <ftp://ftp.isi.edu/internet-drafts/draft-basak-mpls-oxc-issues-01.txt>.

[12] Y Rekhter and Kireeti Kompella. LSP Hierarchy with MPLS TE. Internet Draft (Work in Progress) draft-ietf-mpls-lsp-hierarchy-01, Internet Engineering Task Force, 9 2000. <ftp://ftp.isi.edu/internet-drafts/draft-ietf-mpls-lsp-hierarchy-01.txt>.

[13] Y Rekhter, E. Mannie, John Drake, Debanjan Saha, Don Fedyk, Kireeti Kompella, Vishal Sharma, Greg Bernstein, and Ayan Banerjee. ISIS Extensions in Support of Generalized MPLS. Internet Draft (Work in Progress) draft-ietf-isis-gmpls-extensions-00, Internet Engineering Task Force, 9 2000. <ftp://ftp.isi.edu/internet-drafts/draft-ietf-isis-gmpls-extensions-00.txt>.