

# Technology and Architecture to Enable the Explosive Growth of the Internet

Adel A. M. Saleh, *Fellow, IEEE* and Jane M. Simmons, *Fellow, IEEE*

## ABSTRACT

At current growth rates, Internet traffic will increase by a factor of one thousand in roughly 20 years. It will be challenging for transmission and routing/switching systems to keep pace with this level of growth without requiring prohibitively large increases in network cost and power consumption. We present a high-level vision for addressing these challenges based on both technological and architectural advancements.

## INDEX TERMS

Internet, IP routers, fiber capacity, network evolution, spectral efficiency, traffic grooming, traffic growth

## I. INTRODUCTION

By numerous accounts, Internet traffic continues to exhibit exponential traffic growth. One of the leading sources monitoring Internet traffic levels indicates that the annual rate of growth is currently 40 to 50% [1]. A major IP router vendor has forecast that the Internet will grow at a compound rate of 35% from 2008 to 2013 [2]. As a related benchmark, a large telecommunications carrier has gauged their year-over-year IP traffic growth rate to be 45% [3]. If we assume a compound annual growth rate on the order of 40%, then Internet traffic will increase by a factor of 1,000 in roughly 20 years, as shown in Fig. 1. Attaining this traffic threshold in this relatively short timeframe necessitates developing a high-level network evolution strategy today.

Network capabilities, specifically transmission and routing/switching capabilities, will need to keep pace with the traffic in order for steady Internet growth to be sustained. While the capacity of fiber transmission systems has increased by a factor of roughly 100 over the past decade, attaining another three orders of magnitude growth will be very challenging. Fiber capacity, which once seemed to be almost infinite compared to the traffic requirements, is now approaching its theoretical limit [4]. Furthermore, building large-scale electronic routers and switches is already a challenge, which will only become more of an impediment in the future. Note that simply satisfying the thousandfold growth requirement is not sufficient; it must be done in a manner that is cost effective and power efficient.

A two-pronged approach will likely be needed to meet these challenges: *technological advancements* to increase the realizable capacity of fiber and routers/switches; and *architectural enhancements* that effectively decrease the traffic burden on the network. Section II presents technological and architectural techniques that can meet the thousandfold traffic growth from a transmission perspective. Section III addresses these aspects for IP routers (the focus is on routers as they are more challenging and costly to scale than switches). The analysis considers transmission and routing in the backbone network. As the discussion is looking 20 years into the future, the numbers presented should be interpreted as rough, round estimates as opposed to precise figures. The analysis expands on the work that was originally presented in [5].

---

Adel A. M. Saleh is with DARPA, Arlington, VA. Jane M. Simmons is with Monarch Network Architects, Holmdel, NJ. The views, opinions, and/or findings contained in this article are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the Department of Defense.

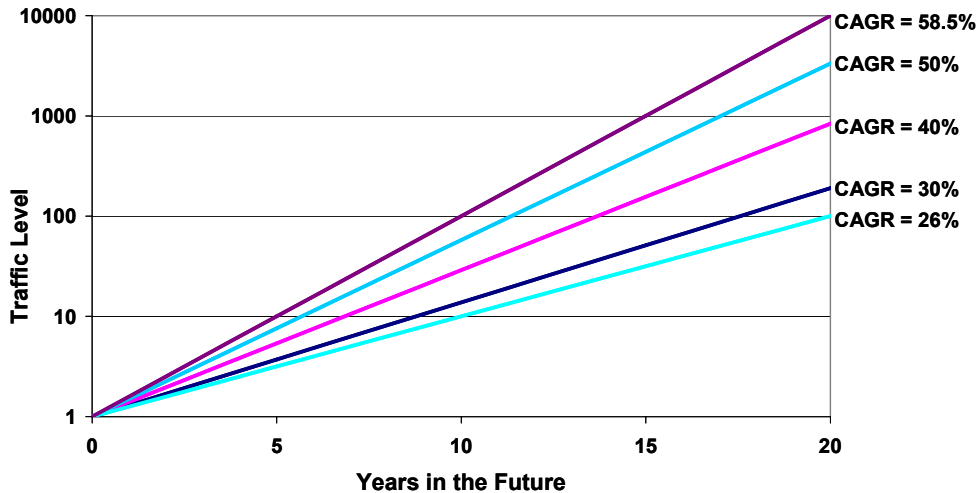


Fig. 1. Level of traffic for various compound annual growth rates (CAGR). At 40% CAGR, the traffic level will increase by a factor of 1,000 in roughly 20 years.

## II. TRANSMISSION

Today’s state-of-the-art deployed technology is capable of supporting 80 wavelengths of 40 Gb/s each, in the C-band. The largest carrier backbone networks currently carry on the order of 4 Tb/s of total traffic. (While we are considering a thousandfold growth in the Internet, we assume that this will manifest itself as a thousandfold growth in the size of carrier networks.) At this level of traffic, and for reasonable assumptions regarding network topology, traffic statistics, and protection implementation, a network equipped with an 80x40 Gb/s transmission system is about one-third full. Thus, to support a thousandfold growth in traffic over today’s levels, a combination of technology and architectural enhancements needs to provide about a factor of 350 benefit with respect to transmission.

Current systems typically support optical bypass, where traffic that is transiting a node can remain in the optical domain, thereby eliminating much of the required electronic terminating equipment. One requirement for optical bypass is an extended optical reach, which is the distance the optical signal can travel before the signal quality degrades to a level that necessitates regeneration. It has been shown that the ‘sweet spot’ for optical reach in continental-scale networks is on the order of 2,000 to 2,500 km [6]. It is highly desirable to maintain an optical reach of this extent as, together with all-optical switching equipment, it has played a significant role over the last decade in reducing network cost and power consumption.

### A. Technology

#### 1. Spectral Efficiency

In the past, transmission systems have kept pace with traffic growth by both increasing the number of wavelengths supported on a fiber and increasing the bit-rate of each wavelength. In the mid to late 1990s, state-of-the-art transmission systems consisted of 16 wavelengths of 2.5 Gb/s each, representing a spectral efficiency of 0.01 bits/s/Hz (spectral efficiency is defined as the ratio of the information bit rate to the total bandwidth consumed). As mentioned above, today’s most advanced deployed systems are 80x40 Gb/s, for a spectral efficiency of 0.8 bits/s/Hz. Additionally, 80x100 Gb/s systems (spectral efficiency of

2.0 bits/s/Hz) are on the horizon. This technological progress has been attained through more complex modulation schemes and more advanced electronic signal processing.

Increased capacity through increased spectral efficiency has provided favorable economies of scale. For example, the 10 Gb/s transponder cost is approximately twice that of a 2.5 Gb/s transponder, resulting in a halving of the cost per bit/s. Similarly, the power consumption and size of a 10 Gb/s transponder is less than that of four 2.5 Gb/s transponders, providing benefits in power and space per bit/s. It is expected that these trajectories will eventually apply to 40 Gb/s and 100 Gb/s equipment as these technologies mature.

Continuing the trend of increased spectral efficiency, however, will become increasingly more difficult. The analysis of [4] indicates that for an optical reach of 2,000 km, the theoretical limit on spectral efficiency is about 6 to 7 bits/s/Hz per polarization. The theoretical limit is unlikely to be attainable in a practical system; thus, it is reasonable to consider the realizable spectral efficiency limit to be more on the order of 4 bits/s/Hz per polarization. If it is assumed that future systems will be dual-polarization (as the planned 100 Gb/s systems are), the realizable system spectral efficiency will likely be on the order of 8 bits/s/Hz. This could be realized, for example, with an 80x400 Gb/s system.

As compared with today's 40 Gb/s systems, 8 bits/s/Hz represents a factor of 10 increase in transmission capacity, which is clearly insufficient to meet long-term projections of traffic growth. Reference [7] also noted the challenge of meeting future traffic growth solely via an increase in spectral efficiency.

## *2. Expanded Transmission Band*

80x40 Gb/s systems are accommodated in approximately 32 nm of spectrum in the C-band. With the increase in spectral efficiency described above, it is expected that 80x400 Gb/s systems would be supported in this band as well. However, expansion into other bands can be used to increase system capacity. For example, the L-band provides low fiber loss comparable to the C-band, making it the most likely choice for expansion. It is important that an expanded system require only a single amplifier across the spectrum, to avoid the cost of multiple band amplifiers. Also, the tunable transponders ideally need to tune across the whole utilized spectrum. We assume that this will be feasible for ~65 nm of spectrum across the C- and L-bands (a current system already supports 54 nm across the C- and L-bands with a single amplifier [8]). Thus, we assume that expanding the transmission band will result in a factor of two increase in system capacity; e.g., a 160x400 Gb/s system.

## *3. Multicore Fiber*

Given that the capacity limit of a single fiber is being approached, it is natural to consider carrying future traffic on multiple fibers. However, this would require that the number of deployed optical amplifiers and the port-size of all-optical switching devices, such as reconfigurable optical add/drop multiplexers (ROADMs), scale by the same amount. Thus, this solution does not provide benefits in cost per bit/s, and equally important, power per bit/s (i.e., energy per bit).

An alternative solution is to increase the number of cores supported in a fiber. While carrier fiber plant is typically composed of single-core, single-mode fiber, there have been recent advances in multicore fiber [9], where ideally the total fiber capacity increases in proportion to the number of cores. In order for the multicore solution to be effective, it must continue to demonstrate the benefits of a conventional single-fiber solution. For example, a single optical amplifier must be capable of amplifying

each of the fiber cores, rather than requiring one amplifier per core. Operationally, a single connector must be capable of interconnecting multicore fibers as opposed to requiring one connector per core.

Multicore fiber presents many challenges, most notably cross-talk between the cores. It will also require that new fiber plant be deployed (although new low-loss fiber plant might be needed to achieve high spectral efficiencies anyway). Additionally, many of the multicore experiments that have been reported cover distances of only about 100 km. The number of cores per fiber in these experiments is typically 7 or 19, as these numbers are compatible with hexagonal packing. For purposes of our discussion, it is assumed that 7 cores will eventually be feasible for long-haul applications. Whether this is realistic given the required optical reach of 2,000 to 2,500 km is unknown at this time; the number of cores may need to be smaller to allow for greater inter-core distance and reduced cross-talk. Clearly, more research is warranted on multicore fiber, as its potential contribution to meeting future capacity requirements is significant. (Another solution being pursued is the use of multimode fiber with electronic multiple-input multiple-output processing, but this is likely to consume a significant amount of power.)

## **B. Architecture**

Increased spectral efficiency, expansion into part of the L-band, and the use of multicore fiber yield about a factor of 140 increase in the capacity of transmission systems (using aggressive assumptions). As this does not meet the target level of growth, networks will need to rely on architectural enhancements to reduce the effective traffic load.

### *1. IP Packing*

IP traffic flows are typically much smaller than the data-rate of a wavelength such that many flows are carried on a single wavelength. Additionally, IP flows have traditionally been quite bursty, demonstrating a high peak-to-average data-rate ratio. To accommodate this burstiness, carriers do not tightly pack IP links (i.e., the IP-carrying wavelengths between routers); they leave ‘headroom’ such that sudden bursts of traffic can be handled without excessive packet loss or delay, and to allow for rerouting under failure conditions. In 2005, it was noted that the average fill-rate of IP links in the US Internet was about 25% [10].

As indicated previously, the data-rate of a wavelength has steadily increased. The average data-rate of individual flows has not increased at the same rate, resulting in a larger number of flows being multiplexed onto a single wavelength. This in turn has had a smoothing effect on the aggregated traffic, lessening the need to drastically overprovision IP links [11, 12]. Using a multi-rate Erlang loss model to determine the maximum utilization that results in an acceptable blocking probability [13] yields utilizations on the order of 95% for 400 Gb/s wavelengths. Depending on the amount of carried best-effort IP traffic that can be scaled back under failure conditions (which affects the amount of overprovisioning needed for rerouting), overall utilizations of 65% or more are feasible. If we assume that the current average fill-rate is on the order of 30 to 35% (representing some improvement over the 2005 figure), then this represents a factor of two benefit. Essentially, roughly twice as much traffic can be packed onto the same number of wavelengths, thereby reducing the effective capacity burden. Of course, traffic that is already at the data-rate of a wavelength (i.e., optical wavelength services) will not realize this benefit. We estimate that no more than 20% of the traffic is likely to be composed of wavelength services, such that the factor of two benefit in capacity applies to at least 80% of the traffic.

It should be noted that IP router developments could also play a role in increasing the utilization level of IP links. For example, “flow routers” are capable of controlling the rate and route of individual flows, such that quality-of-service can be attained through better control rather than through overprovisioning [10].

2. *Multicasting, Asymmetric Traffic, and Improved Caching*

Architectural benefits can also be achieved by better tuning the Internet to the changing nature of its traffic, for example, the growing amount of video distribution. First, we consider multicast as a replacement for multiple unicast connections between a single source and multiple destinations. This is illustrated in Fig. 2, where Node A is transmitting the same payload to Nodes C, F, and G. In Fig. 2(a), three separate connections are established, whereas in Fig. 2(b) a single multicast tree is established. It can readily be seen that the capacity requirements will be smaller with multicast.

To investigate the benefits of multicast, a study was performed on the network shown in Fig. 3. This network represents a typical U.S. backbone network, with 60 nodes and an average nodal degree of 2.6. Five thousand multicast sets were generated, with one source node and  $D$  destination nodes, where  $D$  was uniformly distributed between 5 and 15. The destinations were chosen based on their traffic weightings. (A typical demand set for this network was used to obtain the weightings. However, there was little difference as compared to the case where the nodes were selected with equal likelihood.) The study compared routing  $D$  unicast connections versus one multicast connection. The multicast routing heuristic illustrated in [14] was used to route the multicast trees. The results show that multicast provides a factor of roughly three benefit in capacity, where capacity was measured as the average number of wavelengths required on a link (approximately the same capacity benefits were obtained when capacity was measured in terms of bandwidth-distance, or in terms of the number of wavelengths needed on the most heavily utilized link). While the study specifically investigated multicast at the optical layer, IP multicast offers similar levels of capacity savings (although somewhat smaller due to the finer granularity of the IP layer).

In backbone networks today, connections are almost always bi-directional and symmetric; i.e., if a connection of rate  $R$  is established from Node A to Node Z, then a connection of rate  $R$  is also established from Node Z to Node A. With multicast distribution, the connections only need to be unidirectional, from the source to the destinations. There are also other applications where the traffic may be very asymmetric in the two directions; for example, a 40 Gb/s connection may be needed in one direction, but only a 2.5 Gb/s in the reverse direction. Future networks could take advantage of this asymmetry and only provision what is needed in the two directions, to reduce the amount of utilized capacity. This will necessitate modifying carrier network management systems.

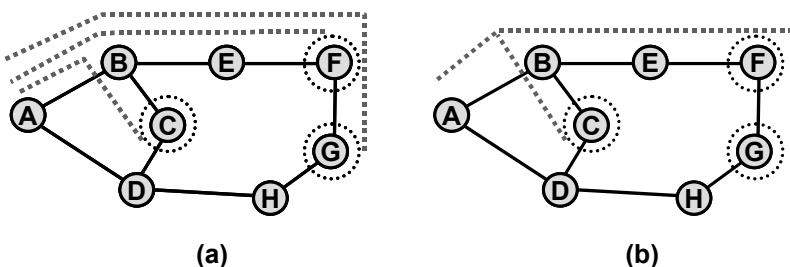


Fig. 2. (a) Three unicast connections from A to C, F, and G. (b) One multicast connection to the same three nodes.

Caching is another architectural feature that can be used to reduce capacity requirements. As the Internet is used more as a repository of data and video, the amount of traffic that can be cached will likely increase. Furthermore, caching algorithms are likely to improve, to increase the probability that the desired data is stored at a nearby location.

While the capacity benefit of multicast was estimated above to be a factor of three, it is more difficult to quantify the gains that can be attained through taking advantage of asymmetric connections and improved caching. We estimate that overall, these factors (including multicast) will produce capacity benefits on the order of a factor of four. Furthermore, we assume that these gains apply to roughly 20% of the traffic.



*Fig. 3. A typical US backbone network, with 60 nodes and an average nodal degree of 2.6.*

### *3. Dynamic Optical Networking*

Networks today are mostly quasi-static, with connection setup typically requiring on-site manual involvement, and connections often remaining established for months or years. As a first move away from this relatively fixed environment, networks are becoming configurable, where connections can be established remotely through software control, assuming the necessary equipment is already deployed in the network. Configurable networks take advantage of flexible network elements such as ROADMs, which can be remotely reconfigured to add, drop, or bypass any wavelength without affecting existing network traffic, and tunable transponders, which can tune to any of the wavelengths supported on a fiber.

The next step in this evolution is dynamic networking, where connections can be rapidly established and torn down, without the involvement of operations personnel. Dynamic offerings are currently limited to sub-wavelength rates with setup times on the order of minutes [15]. However, research is underway to extend dynamism to wavelength services and to provide setup times on the order of 100 msec. to 1 sec. [13]. Dynamic networking takes advantage of advancements such as very fast switching, the ability to manage optical transients, and a distributed control plane. This will likely be a push-pull evolution, where the need for on-demand services (e.g., cloud computing) drives the implementation of dynamic networks, and the availability of a dynamic infrastructure fuels the development of more services that can take advantage of it (e.g., distributed computing on large data sets, interactive visualization and collaboration, etc.).

Dynamic networking is advantageous because it delivers bandwidth where and when it is needed. Bandwidth does not need to be permanently reserved for services that are active for only a small

percentage of the time, in effect, decreasing the network cost and capacity requirements. The benefits for the individual users of dynamic networking can be significant.

A network study on dynamic traffic was performed using the network of Fig. 3. Demands were modeled as on/off services, with an average on-time of 10%. Connections were established and torn down as the demands toggled between the *on* and *off* states. This was compared to a scheme where connections were maintained for the duration of the service regardless of whether the service was active or not. The more traffic that can take advantage of dynamism, the greater the capacity benefits of a dynamic network. For purposes of the study, it was assumed that in a 20-year timeframe, about 25% of the connections in a network would be dynamic (this refers to connections at the optical layer, where a connection can carry a single wavelength service or multiple substrate services). The study showed that with a thousandfold growth in traffic and 25% of the connections carrying on/off traffic, dynamic networking reduces the capacity required *for these services* by a factor of five.

### C. Summary of Transmission Factors

Table 1 summarizes the various factors described above that will either increase the capacity supported on a transmission system or that will decrease the effective traffic load. The third column in the table indicates what percentage of traffic we estimate will be able to take advantage of a particular factor. When combined with the fact that current systems are about one-third full, the technological and architectural advancements will provide transmission support for a thousandfold growth over today’s traffic levels. The aggressive assumptions discussed above are indicative of the challenges faced.

As indicated in the table, the combination of architectural enhancements effectively reduces the total required capacity by a factor of roughly 2.5. This result depends on the assumptions regarding traffic-type percentages. For example, if 50% of the traffic is dynamic rather than 25%, the benefit factor of dynamism increases from 5 to 6, and the combined benefit of architectural enhancements increases from 2.5 to about 3.5.

**Table 1. Summary of the Factors Affecting Transmission**

	<b>Benefit Factor</b>	<b>Percentage of Traffic Subject to Benefit</b>	<b>Effective Capacity Multiplier</b>
<b>Available Excess Capacity in Today’s Networks</b>	<b>3</b>	<b>100%</b>	<b>3</b>
<b>Increased Spectral Efficiency</b>	<b>10</b>	<b>100%</b>	<b>10</b>
<b>Expanded Transmission Band</b>	<b>2</b>	<b>100%</b>	<b>2</b>
<b>Multicore Single-Amp Fiber</b>	<b>7</b>	<b>100%</b>	<b>7</b>
<b>More Efficient IP Packing</b>	<b>2</b>	<b>80%</b>	<b>1.7</b>
<b>Multicast/Asymmetric/Caching</b>	<b>4</b>	<b>20%</b>	<b>1.2</b>
<b>Dynamic Networking</b>	<b>5</b>	<b>25%</b>	<b>1.3</b>
<b>Total Effective Capacity Multiplier</b>			<b>~ 1,100</b>

### **III. IP ROUTERS**

The previous section addressed traffic on a link; this section addresses traffic at a node. Nodal traffic may be handled at multiple network layers. The physical layer of most current networks includes switches, such as ROADMs, that process the traffic on a wavelength granularity in the optical domain. Because the switching is done optically rather than electronically, these elements have demonstrated good scalability in cost, size, and power consumption, and are likely to continue to do so. Such optical networking elements are fundamental to network scalability.

If traffic needs to be processed at a granularity finer than a wavelength, it is sent to an electronic switch or router. Electronic switching or routing can occur at Layer 1 (e.g., in an Optical Transport Network (OTN) switch), Layer 2 (e.g., in an Ethernet switch), or Layer 3 (in an IP router), or some combination of these layers. The exact layered architecture depends on the carrier. In general, the higher the layer, the finer the granularity of operation, and thus, the more challenging it is to scale. Therefore, this section discusses nodal processing in the context of IP routers. If a factor of 1,000 benefit can be achieved at the IP layer, then it is assumed that architectures such as IP-over-OTN-over-WDM, where much of the traffic grooming is offloaded from the IP router to the OTN switch, are feasible as well.

#### **A. Technology**

Today's largest carriers have deployed core IP routers of maximum size about 3 Tb/s. The development of routers of size 160 Tb/s (full duplex) has already been announced [16], which represents more than a factor of 50 increase over today's deployed sizes. It is certainly conceivable that in 20 years, even larger routers will be developed, such that close to a thousandfold increase may be achieved through technology advancements alone, e.g., improved integration density and the use of optical interconnects.

However, the fact that a router of this size can be built does not mean it is practical with respect to cost, size, and power. For example, the power consumption of today's core routers is on the order of 10 Watts per Gb/s [17]. If we assume that in 20 years a 3,000 Tb/s router will exist, then using current power figures, the total power consumption for a single router would be 30 Megawatts. Of course, power improvements can be expected over time; using the estimate of 20% better power efficiency per year [17], over 20 years, yields a single 3,000 Tb/s router consuming 350 Kilowatts.

For our purposes here, it is assumed that in the 20-year timeframe, feasible routers will be about five times larger than what has already been announced; e.g., ~750 Tb/s, which represents a factor of 250 increase over today's deployed sizes. Even if larger routers are possible, the operational challenges of deploying such a large device are impetus to consider architectural innovations that can mitigate their need.

#### **B. Architecture**

##### *1. IP Packing*

As discussed in Section II.B.1, increased wavelength data-rates result in less bursty aggregated substrate traffic, allowing for higher fill-rates. It was estimated that the average IP-link fill-rate would increase by a factor of two, such that the number of required wavelengths to carry the substrate services would decrease by approximately the same factor. This translates into half as many ports needed on the IP routers (it is assumed that wavelength services do not enter the routers). Thus, tighter IP packing will result in approximately a factor of two decrease in the required size of IP routers.

## 2. Optical Aggregation at the Network Edge

As the level of network traffic grows while the number of network nodes remains approximately fixed, the average amount of traffic between node pairs increases. Thus, an increasing amount of traffic can be efficiently packed into wavelengths at the edge of the network without requiring further packing in the backbone network (given that the wavelength data-rate does not increase in proportion to the traffic). This implies that efficient packing in the edge networks (i.e., regional and metro-core networks) can be used to offload much of the burden of the core IP routers, a paradigm that was discussed in [18] and illustrated in Fig. 4. (This paradigm is also similar to optical flow switching [19].) Optical aggregation techniques, such as optical burst switching, can possibly be used at the edge to reduce the overall amount of electronic processing. Optical aggregation is typically more suitable for edge networks than for backbone networks, due to these techniques requiring collision management and/or scheduling.

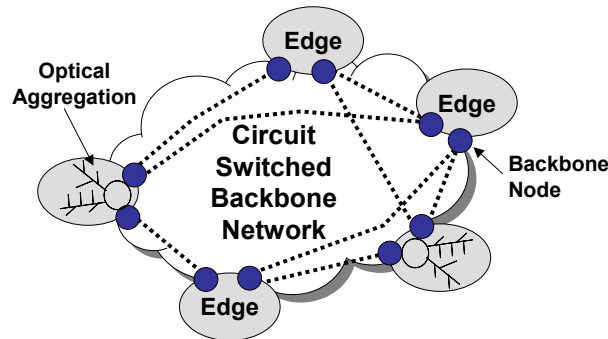


Fig. 4. Optical aggregation is implemented in the edge network, delivering traffic to the backbone nodes; this traffic does not undergo further grooming in the backbone network. The edge/backbone interface could be all-optical or optical-electronic-optical as discussed in more detail in [18].

To explore the benefits of optical edge aggregation, a study was performed on the network of Fig. 3, with the level of traffic assumed in this paper combined with 400 Gb/s wavelengths. It is unlikely that all traffic could take advantage of optical aggregation at the edge as this will require specialized optical equipment, which may not be installed in a particular metro network or may be too costly for some end customers. Thus, it was assumed that 50% of the substrate traffic was eligible to be optically aggregated at the edge, with the remaining 50% being processed by IP routers as usual. This resulted in the maximum required size of the backbone IP router decreasing by a factor of two as compared to a conventional architecture. (Reference [18] aggressively assumed 90% of the traffic was optically aggregated at the edge, yielding a factor of ten benefit.) Since the optically aggregated traffic does not take advantage of grooming inside the backbone network, the wavelengths are slightly less packed, but the overall capacity requirement increase is negligible.

Another approach to reducing the burden on IP routers is to support a range of wavelength data-rates on a fiber, for example, from 40 Gb/s to 400 Gb/s (where the 40 Gb/s wavelengths are more tightly spaced to maintain the same spectral efficiency). By better matching the wavelength data-rate to the service data-rate, less traffic grooming would be needed in the IP routers, as is discussed in [20]. However, supporting closely spaced wavelengths (e.g., 40 Gb/s wavelengths at 5 GHz spacing) will be challenging. While we do not include this technique in our assumptions here, it may be worthy of future research.

### 3. Multicasting, Asymmetric Traffic, Improved Caching, and Dynamic Optical Networking

Similar to Sections II.B.2 and II.B.3, taking advantage of multicasting, asymmetric traffic, improved caching, and dynamic optical networking will reduce routing requirements as well. However, some of the connections that will benefit from these factors represent wavelength services, which we assume bypass the IP routers. Furthermore, we assume that optical edge aggregation is implemented, as discussed in the previous section, such that much of the substrate traffic that will benefit from these factors has already been offloaded from the IP routers. With these assumptions, the contribution to router-size reduction from these factors will be small, and hence it is not included in our estimates.

### C. Summary of IP Router Factors

Table 2 summarizes the various factors discussed above that relate to IP router size. Overall, the total benefit is 1,000, which meets the target traffic growth level.

**Table 2. Summary of the Factors Affecting Maximum Required IP Router Size**

	<b>Benefit Factor</b>	<b>Percentage of Router Capacity Affected by Factor</b>	<b>Effective Capacity Multiplier</b>
<b>Larger IP Routers</b>	<b>250</b>	<b>100%</b>	<b>250</b>
<b>More Efficient IP Packing</b>	<b>2</b>	<b>100%</b>	<b>2</b>
<b>Optical Edge Aggregation</b>	<b>2</b>	<b>100%</b>	<b>2</b>
<b>Total Effective Capacity Multiplier</b>			<b>1,000</b>

Note that this discussion on routers has implicitly assumed that today’s networks implement optical bypass, such that IP traffic does not have to enter a router at every node along its path. This is already yielding a factor of two reduction in the maximum required size of a router in today’s large networks [14].

### IV. CONCLUSION

At the current pace of growth, Internet traffic is doubling approximately every two years, leading to a factor of 1,000 growth in the next two decades. We have shown that such staggering growth can indeed be supported, while keeping the network cost and power consumption in check. This requires advances in both technology, to increase the capacity of transmission and routing/switching systems, and architecture, to effectively reduce the capacity requirements. While we have addressed only the backbone portion of the network, access networks will need to scale as well, through a combination of advanced broadband fiber, cable, and wireless technologies. Of course, while the pace of Internet growth can be expected to slow at some point, eventually the thousandfold growth figure will be exceeded as well, requiring even further innovations, yet to be invented!

## References

- [1] A. Odlyzko, Minnesota Internet Traffic Studies (MINTS), Available: [www.dtc.umn.edu/mints/home.html](http://www.dtc.umn.edu/mints/home.html).
- [2] Cisco Visual Networking Index: Forecast and Methodology, 2008–2013, White Paper, June 9, 2009.
- [3] A. Colby, “AT&T, NEC, Corning complete record-breaking fiber capacity test,” May 11, 2009, Available: [news.soft32.com/att-nec-corning-complete-record-breaking-fiber-capacity-test\\_7372.html](http://news.soft32.com/att-nec-corning-complete-record-breaking-fiber-capacity-test_7372.html).
- [4] R.-J. Essiambre, et al., “Capacity limits of optical fiber networks,” *Journal of Lightwave Technology*, vol. 28, no. 4, February 15, 2010, pp. 662-701.
- [5] A. A. M. Saleh, “Dynamic optical networking to enable scalability of the future Internet,” *OFC/NFOEC Future Internet Symposium*, San Diego, California, Feb. 24-28, 2008, Available: [www.monarchna.com/FutureInternet-OFC-NFOEC-2008-Saleh.pdf](http://www.monarchna.com/FutureInternet-OFC-NFOEC-2008-Saleh.pdf).
- [6] J. M. Simmons, “On determining the optimal optical reach for a long-haul network,” *Journal of Lightwave Technology*, vol. 23, no. 3, Mar. 2005, pp. 1039-1048.
- [7] A. R. Chraplyvy, “The coming capacity crunch,” *ECOC Plenary Talk*, Vienna, Austria, Sept. 20-24, 2009, Paper 1.0.2.
- [8] D. A. Fishman, W. A. Thompson, and L. Vallone, “LambdaXtreme® Transport System: R&D of a high capacity system for low cost, ultra long haul DWDM transport,” *Bell Labs Technical Journal*, vol. 11, no. 2, Summer 2006, pp. 27–53.
- [9] K. Imamura, et al., “Multi-core holey fibers for the long-distance (>>100 km) ultra large capacity transmission,” *OFC 2009*, San Diego, CA, March 22-26, 2009.
- [10] L. Roberts, “Enabling data-intensive iGrid applications with advanced network technology,” *iGrid 2005*, San Diego, CA, Sept. 26-29, 2005.
- [11] L. Yao, M. Agapie, J. Ganbar, and M. Doroslovacki, “Long range dependence in internet backbone traffic,” *ICC 2003*, Anchorage, AK, May 11-15, 2003.
- [12] Ciena, “Evolution to the 100G transport network,” White Paper, Nov. 2007.
- [13] A. Chiu, et al., “Network design and architectures for highly dynamic next-generation IP-over-optical long distance networks,” *Journal of Lightwave Technology*, vol. 27, no. 12, June 15, 2009, pp. 1878-1890.
- [14] J. M. Simmons, *Optical Network Design and Planning*, Springer, New York, 2008.
- [15] AT&T Product Brief, “AT&T Optical Mesh Service – OMS,” May 13, 2008, Available: [www.business.att.com/content/productbrochures/PB-OMS\\_16312\\_V01\\_05-13.pdf](http://www.business.att.com/content/productbrochures/PB-OMS_16312_V01_05-13.pdf).
- [16] Cisco, “Advanced platform designed to deliver new wave of video, mobile and data center/cloud services,” Press Release, March 9, 2010.
- [17] J. Baliga, et al., “Energy consumption in optical IP networks,” *Journal Of Lightwave Technology*, vol. 27, no. 13, July 1, 2009, pp. 2391-2403.
- [18] A. A. M. Saleh and J. M. Simmons, “Evolution toward the next-generation core optical network,” *Journal of Lightwave Technology*, vol. 24, no. 9, Sept. 2006, pp. 3303-3321.
- [19] G. Weichenberg, V. Chan, and M. Medard, “Design and analysis of optical flow-switched networks,” *IEEE/OSA Journal of Optical Communications and Networking*, vol. 1, no. 3, Aug. 2009, pp. B81 - B97.
- [20] J. Berthold, A. A. M. Saleh, L. Blair, and J. M. Simmons, “Optical networking: past, present, and future,” *Journal Of Lightwave Technology*, vol. 26, no. 9, May 1, 2008, pp. 1104-1118.