Is IP going to take over the world (of communications)?

Pablo Molinero-Fernández
Stanford University
molinero@stanford.edu

Nick McKeown
Stanford University
nickm@stanford.edu

Hui Zhang
Turin Networks and Carnegie Mellon University
hzhang+@cs.cmu.edu

ABSTRACT

While it is technically pleasing to believe that IP will dominate all forms of communication, our delight in its elegance is making us overlook its shortcomings. IP is an excellent means to exchange data, which explains its success. It remains ill suited as a means to provide many other types of service; and is too crude to form the transport infrastructure in its own right. To allow the continued success of IP, we must be open-minded to it living alongside, and co-operating with other techniques (such as circuit switching) and protocols that are optimized to different needs. In this position paper, we question some of the folklore surrounding IP and packet switching. We conclude that while packet-switched IP will continue to dominate best-effort data services at the edge of the network, the core of the network will use optical circuit switching as a platform for multiple services.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Circuit-switching networks, Packet-switching networks*; C.2.5 [**Computer-Communication Networks**]: Local and Wide-Area Networks—*Internet (e.g., TCP/IP)*

Keywords

IP, packet switching, circuit switching

1. INTRODUCTION

Whatever the initial goals of the Internet, there are two main characteristics that seem to account for its success: *reachability* and *heterogeneity*. IP provides a simple, single, global address to reach every host, enables unfettered access between all hosts, and adapts the topology to restore reachability when links and routers fail. IP hides heterogeneity in the sense that it provides a single, simple service abstraction that is largely independent of the physical links over which it runs. As a result, IP provides service to a huge variety of applications and operates over extremely diverse link technologies.

The growth and success of IP has given rise to some widely held assumptions amongst researchers, the networking industry and the public at large. One common assumption is that it is only a matter of time before IP becomes the sole global communication infrastructure, dwarfing and eventually displacing existing communication infrastructures such as telephone, cable and TV networks. IP is already universally used for data networking in wired networks (enterprise networks and public Internet), and is being rapidly adopted for data communications in wireless and mobile networks. IP is increasingly used for both local and long-distance voice communications, and it is technically feasible for packet-switched IP to replace SONET/SDH.

A related assumption is that IP Routers (based on packet-switching) will become the most important, or perhaps only, type of switching device inside the network. This is based on our collective belief that packet switching is inherently superior to circuit switching because of the efficiencies of statistical multiplexing, and the ability of IP to route around failures. It is widely assumed that IP is simpler than circuit switching, and should be more economical to deploy and manage. And with continued advances in the underlying technology, we will no doubt see faster and faster links and routers.

On the face of it, these assumptions are quite reasonable. Technically, IP is flexible enough to support all communication needs, from best-effort to real-time. With robust enough routers and routing protocols, and with extensions such as weighted fair queueing, it is possible to build a packet-switched, datagram network that can support any type of application, regardless of their requirements.

But for all its strengths, we (the authors) do not believe that IP will displace existing networks; in fact, we believe that many of the assumptions discussed above are not supported by reality, and do not stand up to close scrutiny.

It is the goal of this paper to question the assumption that IP will be *the* network of the future. We will conclude that if we started over - with a clean slate - it is not clear that we would argue for a universal, packet-switched IP network. We believe that in the future, more and more users and applications will demand predictability from the Internet;

both in terms of the availability of service, and the timely delivery of data. IP was not optimized to provide either, and so it seems unlikely to displace networks that already provide both.

We take the position that while IP will be the network layer of choice for best-effort, non-mission critical and non-real-time data communications (such as information exchange and retrieval), it will live alongside other networks, such as circuit-switched networks, that are optimized for high revenue time-sensitive applications that demand timely delivery of data and guaranteed availability of service.

We realize that our position is a controversial one. But regardless of whether or not we are correct, as researchers we need to be prepared to take a step back, to take a hard look at the pros and cons of IP, and its likely future. As a research and education community, we need to start thinking how IP will co-exist and co-operate with other networking technologies.

2. IP FOLKLORE

In what follows, we try to identify some folkloric assumptions about IP and the Internet, and examine each in turn. We will start with the most basic assumption, and easiest to dispel: that the Internet *already* dominates global communications. This is not true by any reasonable metric: market size, number of users, or the amount of traffic. Of course, the Internet has not yet reached maturity, and it may still grow to dominate the global communications infrastructure. We should ask ourselves if packet-switched IP offers inherent and compelling advantages that will lead to its inevitable and unavoidable dominance. This requires us to examine some "sacred cows" of networking; for example, that packet switching is more efficient than circuit switching, that IP is simpler, it lowers the cost of ownership, and it is more robust.

2.1 IP already dominates global communications

Although the Internet has been a phenomenal success, it is currently only a small fraction of the global communication infrastructure consisting of separate networks for telephones, broadcast TV, cable TV, satellite, radio, public and private data networks, and the Internet. In terms of revenue, the Internet is a relatively small business. The US business and consumer-oriented ISP markets have revenues of $13B each (2000) [5] [6], by contrast, the TV broadcast industry has revenues of $29.8B (1997), the cable distribution industry $35.0B (1997), the radio broadcast industry $10.6B (1997) [31], and the phone industry $268.5B (1999), of which $111.3B correspond to long distance and $48.5B to wireless [13]. The Internet reaches 59% of US households [22], compared to 94% for telephones and 98% for TV [20, 25]. It is interesting to note that, if the revenue per household remains the same, the total revenue for the ISP industry can at most double.

If we restrict our focus to the data and telephony infrastructure, the core IP router market still represents a small fraction of the public infrastructure, contrary to what happens in the private enterprise data networks. As shown in Table 1 the expenditure on core routers worldwide was $1.7B in 2001, compared to $28.0B for transport circuit switches. So in terms of market size, revenue, number of users, and expenditure on infrastructure, it is safe to say that the Internet does not currently dominate the global communications infrastructure.

The current infrastructure consists of a transport network - made of circuit-switched SONET and DWDM devices - on top of which run multiple service networks. The service networks include the voice network (circuit switched), the IP network (datagram, packet switched), and the ATM/Frame Relay networks (virtual-circuit switched). When considering whether IP has or will take over the world of communications, we need to consider both the transport and service layers.

In what follows, we will be examining which of two outcomes is more likely: Will the packet-switched IP network grow to dominate and displace the circuit switched transport network; or will the (enhanced) circuit-switched TDM and optical switches continue to dominate the core transport network?

Segment	Market size
Core routers	$1.7B
Edge routers	$2.4B
SONET/SDH/WDM	$28.0B
Telecom MSS	$4.5B

Table 1: World market breakup for the public telecommunications infrastructure in 2001 [30].

2.2 IP is more efficient

> "Analysts say [packet switched networks] can carry 6 to 10 times the traffic of traditional circuit-switched networks" – **Business Week**.

From the early days of computer networking, it has been well known that packet switching makes efficient use of scarce link bandwidth [1]. With packet switching, statistical multiplexing allows link bandwidth to be shared by all users, and work-conserving link sharing policies (such as FCFS and WFQ) ensure that a link is always busy when packets are queued-up waiting to use it. By contrast, with circuit switching, each flow is assigned to its own channel, so a channel could go idle even if other flows are waiting. Packet switching (and thus IP) makes more efficient use of the bandwidth than circuit switching, which was particularly important in the early days of the Internet when long haul links were slow, congested and expensive.

It is worth asking: What is the current utilization of the Internet, and how much does efficiency matter today? Odlyzko and Coffman [24, 9] report that the average link utilization in links in the core of the Internet is between 3% and 20% (compared to 33% average link utilization in long-distance phone lines [24, 29]). The reasons that they give for low utilization are threefold; First, Internet traffic is extremely asymmetric and bursty, but links are symmetric and of fixed capacity, second it is difficult to predict traffic growth in a link, so operators tend to add bandwidth aggressively, and finally as faster technology appears it is more economical to add capacity in large increments.

There are other reasons to keep network utilization low. When congested, a packet-switched network performs badly, becomes unstable, and can experience oscillations and synchronization. Because routing protocol packets are transmitted in-band, they can be lost or delayed due to network congestion or control processor overload. This causes inconsistent routing state, and may result in traffic loops, black holes and disconnected regions of the network, which further exacerbates congestion in the data path [17]. Today, network providers address these problems by keeping network utilization low.

But perhaps the biggest reason that network providers over-provision their network is to give low packet delay. Users want predictable behavior, which means low queueing delay, even under abnormal conditions (such as the failure of several links and routers). As users, we already demand (and are willing to pay for) huge over-provisioning of Ethernet networks (the average utilization of an Ethernet network today is about 1% [9]) just so we do not have to share the network with others, and so that our packets can pass through without queueing delay. We will demand the same behavior from the Internet as a whole. We will pay network providers to stop using statistical multiplexing, and to instead over-provision their networks, as if it were circuit switched [12]. The demand for lower delay will drive providers to decrease link utilization even lower than it is today.

But simply reducing the *average* link utilization will not be enough to make users happy. For a typical user to experience low utilization, the *variance* of the network utilization needs to be low, too. Reducing variations in link utilization is hard; today we lack effective techniques to do it. It might be argued that the problem will be solved by research efforts on traffic management, congestion control, and multipath routing. But to-date, despite these problems being understood for many years, effective measures are yet to be introduced.

On a related note, we might ask whether users experience lower delay in a packet switched or circuit switched network. Intuition suggests that packet switching will lead to lower delay: A packet switched network easily supports heterogeneous flow rates, and flows can always make forward progress because of processor-sharing in the routers. In practice, we find that it doesn't make much difference whether we use packet switching or circuit switching. We explored this in detail in some earlier work, where we studied (using analysis and simulation) the effect of replacing the core of the network with dynamic fine-granularity circuit switches [19]. Let's define the user response time to be the time from when a user requests a file, until this file finishes downloading. Web browsing and file sharing represent over 65% of Internet transferred bytes today [7], and so the request/response model is representative of typical user behavior. Now consider two types of network: One is the current packet-switched network in which packets share links. In the other network each new application flow triggers the creation of a low bandwidth circuit in the core of the network, similar to what happens in the phone network. If there are no circuits available, the flow is blocked until a channel is free. At the core of the network, where the rate of a single flow is limited by the data-rate of its access link, simulations and analysis in [19] indicate that the user response time is essentially the same for packet switching and circuit switching, independent of the flow length distribution.

In summary, we have observed that packet switching can lead to more efficient link utilization. While efficiency was once a critical factor, it is so outweighed by our need for predictability, stability, immediate access and low delay that network operators are forced to run their networks at a low utilization, forfeiting the benefits of statistical multiplexing.

2.3 IP is robust

> "The Internet was born during the cold war 30 years ago. The US Department of Defence [decided] to explore the possibility of a communication network that could survive a nuclear attack." – **BBC**

The Internet was designed to withstand a catastrophic event where a large number of links and routers were destroyed. This goal is in line with users and businesses who rely more and more on the network connectivity for their activities and operations, and who want the network to be available at all times. Much has been claimed about the reliability of the current Internet, and it is widely believed to be inherently more robust. Its robustness comes from using soft-state routing information; upon a link or router failure it can quickly update the routing tables and direct packets around the failed element.

The reliability of the current Internet has been studied by Labovitz et al. [17]. They have studied different ISPs over several months, and report a median network availability equivalent to a downtime of 471 min/year. By contrast Kuhn [15] found that the average downtime in phone networks is less than 5 min/year. As users we have all experienced network downtime when our link is unavailable, or some part of the network is unreachable. On occasions, connectivity is lost for long periods while routers reconfigure their tables and converge to a new topology. Labovitz et al. [16] observed that the Internet recovers slowly, with a median BGP convergence time of 3 minutes, and frequently taking over 15 minutes. By contrast, SONET/SDH rings, through the use of pre-computed backup paths, are required to recover in less than 50 ms; a glitch that is barely noticeable by the user.

As discussed in Section 2.2, the likelihood of a network getting into a inconsistent routing state is much higher in IP networks because (a) the routing packets are transmitted in-band, and therefore are more likely to incur congestion due to high load of user traffic; (b) the routing computation in IP networks is very complex, therefore, it is more likely for the control processor to be overloaded; (c) the probability of mis-configuring a router is high. And mis-configuration of even a single router may cause instability of a large portion of the network. It is surprising is that we have continued to use routing protocols that allow one badly behaved router to make the whole network inoperable [18]. In contrast, high availability has always been a government-mandated requirement for the telephone network, and so steps have been taken to ensure that it is an extremely robust infrastructure. In circuit networks control messages are usually transmitted over a separate channel or network, and the routing is much simpler.

In datagram networks, inconsistence routing state may cause black holes or traffic loops so that the service to existing user traffic is disrupted – i.e. inconsistent routing is *service impacting*. In circuit networks, inconsistent routing state may result in unnecessary rejection of request for new circuits, but none of the established circuits is affected. In summary, currently with IP, not only are failures more common, but they take longer to be repaired and their impact is deeper.

The key point here is that there is nothing inherently unreliable about circuit switching, and we have proof that it is both possible and economically viable to build a robust circuit-switched infrastructure, that is able to quickly reconfigure around failures. There is no evidence yet that we can define and implement the dynamic routing protocols to make the packet-switched Internet as robust. Perhaps the problems with BGP will be fixed over time and the Internet will become more reliable. But it is a mistake to believe that packet switching is inherently more robust. In fact, the opposite may be true.

2.4 IP is simpler

> "IP-only networks are much easier and simpler to manage, leading to improved economics." – **Business Communications Review**

It is an oft-stated principle of the Internet that the complexity belongs at the end-points, so as to keep the routers simple and streamlined. While the general abstraction and protocol specification are simple, implementing a high performance router and operating an IP network are extremely challenging tasks, particularly as the line rates increase.

If we are looking for simplicity, we can do well to look at how circuit-switched transport switches are built. First, the software is simpler. The software running in a typical transport switch is based on about three million lines of source code [28], whereas Cisco's Internet Operating System (IOS) is based on eight million [10], over twice as many. Routers have a reputation for being unreliable, crashing frequently and taking a long time to restart. So much so that router vendors frequently compete on the reliability of their software.

The hardware in the forwarding path of a circuit switch is also simpler than that of a router. At the very least, the line card of a router must unframe/frame the packet, process its header, find the longest-matching prefix that matches the destination address, decrement the TTL, process optional headers, and then buffer the packet. If multiple service levels are added (e.g., DiffServ [3]), then multiple queues must be maintained, as well as an output link scheduling mechanism.

On the other hand, the linecard of an electronic transport switch typically contains a SONET framer to interface to the external line, a chip to map ingress time slots to egress time slots, and an interface to a switch fabric. Essentially, one can build a transport linecard [27] by starting with a router linecard [26] and then removing most of the functionality.

One measure of this complexity is the number of logic gates implemented in the linecard of a router. An OC192c POS linecard today contains about 30 million gates in ASICs, plus at least one CPU, 300Mbytes of packet buffers, 2Mbytes of forwarding table, and 10Mbytes of other state memory. The trend in routers has been to put more and more functionality on the forwarding path: first, support for multicast (which is rarely used), and now support for QoS, access control, security and VPNs (and we thought that all the complexity was in the end system!). By contrast, the linecard of a typical transport switch contains a quarter of the number of gates, no CPU, no packet buffer, no forwarding table, and an on-chip state memory (included in the gate count). Because they use simpler hardware, electronic circuit switches consume less power, allowing more capacity to be placed in a single rack. It should come as no surprise that the highest capacity commercial transport switches have two to twelve times the capacity of an IP router [8, 21, 14], and sell for about half to 1/12 per gigabit per second. So even if packet switching might be simpler for low data rates, it becomes more complex for high data rates. IP's "simplicity" does not scale.

It is interesting to explore how optical technology will affect the performance of routers and circuit switches. In recent years, there has been a lot of discussion about all-optical Internet routers. There are two reasons why this does not make sense. First, a router is a packet switch, and so inherently requires large buffers to hold packets during times of congestion, and there are currently no economically feasible ways to buffer large numbers of packets optically. The second reason is that an Internet router must perform an address lookup for each arriving packet. Neither the size of the routing table, nor the nature of the lookup, lends itself to implementation using optics.

Optical switching technology is much better suited to circuit switches. Devices such as tunable lasers, MEMS switches, fiber amplifiers and DWDM multiplexers provide the technology to build extremely high capacity, low power circuit switches that are well beyond the capacities possible in electronic routers [2].

2.5 Support of telephony and other real-time applications over IP networks

> "All critical elements now exist for implementing a QoS-enabled IP network." – **IEEE Communications Magazine**

There is a widely-held assumption that IP network can support telephony and other real-time applications. If we look more closely, we find that the reasons for such an optimistic assumption are quite diverse. One school holds the view that IP is ready today: IP networks are and will continue to be heavily over-provisioned, and the average packet delay in the network will be low enough to satisfy the real-time requirements of these applications. These real-time applications, including telephony, can tolerate occasional packet delay/loss and *adapt* to these network variabilities. While today's IP networks are heavily over-provisioned, it is doubtful whether a new solution (far from complete yet) that provides a worse performance can displace the reliable and high quality of service (QoS) provided by today's TDM-based infrastructure (which is already paid-for).

Another school believes that for IP to succeed, it is critical for IP to provide QoS with the same guarantees as TDM but with more flexibility. In addition, the belief is that there is no fundamental technical barrier to build a connection-oriented service (Tenet [11] and IntServ [4]) and to provide guaranteed services in the Internet. Unfortunately, after more than 10 years of extensive research and efforts in the standards bodies, the prospect of end-to-end per-flow QoS in the Internet is nowhere in sight. The difficulty seems to be the huge culture gap between the connection and datagram design communities. By blaming the failure on "connections", a third school holds the view that a simpler QoS mechanism such as DiffServ is the right way to go. Again, we are several years into the process, and it is not at all clear that the "fuzzy" QoS provided by DiffServ will be good enough for customers who are used to the simple QoS provided by the existing circuit-switched transport networks.

Finally, no matter what technology we intend to use to carry voice over the Internet, there are few financial incentives to do so. As Mike O'Dell[1] recently said [23]: "[to have a Voice-over-IP service network one has to] create the most expensive data service to run an application for which people are willing to pay less money everyday, [...] and for which telephony already provides a better solution with a marginal cost of almost zero."

3. DISCUSSION

Up until this point, we have considered some of the folklore surrounding the packet-switched Internet. Our overall goal is to provoke discussion and research on fundamental issues that need to be addressed so that IP can continue to revolutionize the world of communications. We hope to provide a vantage point for the IP community to reflect upon the problems that still need to be solved.

3.1 Dependability of IP networks

High dependability, in the broadest sense, is a must if IP is to become the universal infrastructure for high value applications. For example, voice services and private lines are a high-revenue, and very profitable business. Trusting them to today's unreliable, and unpredictable IP networks would be an unnecessary risk, which is why — despite predictions to the contrary — telephone carriers have not done so.

High dependability means several things: robustness and stability, traffic isolation, traffic engineering, fault isolation, manageability, and last but not least, the ability to provide predictable performance in terms of bounded delay and guaranteed bandwidth (QoS). In its current form, IP excels in none of these areas. Although it is clearly a challenge to achieve each of these goals, they must all be solved for IP to become dependable enough to be used as a transport mechanism.

[1]former Senior Vice President of UUNET, responsible for technical strategic direction and architecture of the network.

3.2 How IP should interact with circuits

The current Internet is based on packet switched routers, interconnected by a circuit switched transport network. Given the benefits of circuit switching, it is inconceivable to us that the network providers would remove the existing, robust, reliable, predictable and largely paid-for transport network, and replace it with a technology that seems more complex, less reliable, more expensive and not yet installed.

What seems more likely is that packet switching will continue to exist at the edge of the network, aggregating and multiplexing traffic from heterogeneous sources for applications that have no delay or quality requirements.

At the core of the network, we expect the circuit switched transport network to remain as a means to interconnect the packet switched routers, and as a means to provide high reliability, and performance guarantees. Over time, more and more optical technology will be introduced into the transport network, leading to capacities that electronic routers cannot achieve.

However, IP will not be the only way to access the circuit switched transport network. Because the packet switched network is unlikely to provide the predictability needed for voice traffic, voice will continue to operate over its own, separate circuit switched edge network and be carried over the shared transport network at the core. This leads us to believe that it is more likely that the routers will be allocated a significant fraction of the circuit switched transport infrastructure, which they can control and adapt to best serve their needs. Such a system has the benefit of enabling IP to gain the benefits of fast optical circuit switches in the core, yet maintain the simple service model for heterogeneous sources at the edge.

3.3 What if we started with a clean slate

In the preceding discussion, we predicted an outcome based on historical reasons, in the context of a pre-existing circuit switched transport network. So if we started again, with the benefit of hindsight, would we build a network with circuit switching at the core, and packet switching at the edge? We believe that we would, and that it would look something like this:

Switching in the edges of the network. Packet switching would be used in the edges of the network as well as those links where bandwidth is scarce (such as wireless access links, satellite links, and underwater cables). The reasons for this are twofold. First, packet switching makes a very efficient use of the bandwidth in these cases. Second, it can greatly improve the end-user response time by borrowing all available link bandwidth when other users are not active. The packet-switched network should ideally gather traffic from disparate sources, and multiplex it together in preparation for carriage over a very high capacity, central, circuit-switched core.

Switching in the core of the network. At the core of the network, there seem to be a number of compelling reasons to use circuit switching: Circuit switching has already demonstrated its robustness, and its ability to quickly recover from failures. It is inherently simpler than packet switching, requiring less work to forward data, and so will cost less per capacity unit as a result, will consume less power, and will take up less space. Last, though probably first, circuit switching provides an easy way to adopt the huge potential of high capacity optical switches at low cost.

Integration of both switching mechanisms. Rather than working independently, both of these mechanisms should be tightly integrated, in such a way that an action in one provokes a reaction in the other. For example, packet switching would have to export the QoS and connection oriented nature of the circuit switched core to the applications that require it. Similarly circuit switching has to respond to the increases in traffic of packet switching, by adapting its capacity among core/edge gateways accordingly. Additionally, we will find more hybrid switches that can do both circuit and packet switching, serving as gateways between the two worlds.

4. REFERENCES

[1] P. Baran. Introduction to distributed communications networks. Memorandum RM-3420-PR, Rand Corporation, Aug 1964.

[2] D. J. Bishop, C. R. Giles, and G. P. Austin. The Lucent LambdaRouter: MEMS technology of the future here today. *IEEE Communications Magazine*, 40(3):75–79, Mar 2002.

[3] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. RFC 2475: An architecture for differentiated services, Dec. 1998.

[4] R. Braden, D. Clark, and S. Shenker. RFC 1633: Integrated services in the Internet architecture: an overview, June 1994.

[5] Cahners. 2001 business isps: Service, size, and share. advanced carrier business report. Report, Cahners, Oct 2001.

[6] Cahners. Information alert newsletter. volume #23. Newsletter, Cahners, July 2001.

[7] CAIDA, Cooperative Association for Internet Data Analysis. *OC48 analysis summary: Distributions of traffic stratified by application*, 2002.

[8] Ciena. *CIENA MultiWave CoreDirector*, 2001.

[9] K. Coffman and A. Odlyzko. *Handbook of Massive Data Sets*, chapter Internet growth: Is there a "Moore's Law" for data traffic? J. Abello, P. M. Pardalos, and M. G. C. Resende editors, Kluwer, 2001.

[10] C. Edwards. Panel weighs hardware, software design options. *EE Times*, June 2000.

[11] D. Ferrari. Real-time communication in an internetwork. *Journal of High Speed Networks IOS Press*, 1(1):79–103, 1992.

[12] C. Fraleigh. *Provisioning IP Backbone Networks to Support Delay Sensitive Traffic*. PhD thesis, Electrical Engineering Dept., Stanford University, 2002.

[13] Industry Analysis Division, Common Carriers Bureau. Trends in telephone service. Report, US Federal Communications Commission, Aug 2001.

[14] Juniper Networks. *T640 Internet Routing Node: Datasheet*, 2002.

[15] R. Kuhn. Sources of failure in the public switched telephone network. *IEEE Computer*, 30(4):31–36, April 1997.

[16] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed internet routing convergence. *IEEE/ACM Transactions On Networking*, 9(3):293–306, June 2001.

[17] C. Labovitz, A. Ahuja, and F. Jahanian. Experimental study of internet stability and wide-area backbone failures. In *Proceedings of FTCS*, Madison, WI, June 1999.

[18] C. Labovitz, R. Wattenhofer, S. Venkatachary, and A. Ahuja. Resilience characteristics of the internet backbone routing infrastructure. In *Proceedings of the Third Information Survivability Workshop*, Boston, MA, October 2000.

[19] P. Molinero-Fernández and N. McKeown. TCP Switching: Exposing circuits to IP. *IEEE Micro Magazine*, 22(1):82–89, Jan/Feb 2002.

[20] National Telecommunications and Information Administration. Falling through the net: Defining the digital divide. Technical report, US Department of Commerce, 1999.

[21] Nortel Networks. *OPTera Connect HDX optical switch*, 2002.

[22] Nua Internet Surveys. *How Many On-line?*, April 2002.

[23] M. O'Dell. Keynote speech. Sigcomm conference, August 2002.

[24] A. Odlyzko. Data networks are mostly empty and for good reason. *IT Professional*, 1(2):67–69, Mar/Apr 1999.

[25] A. Penenberg. The war for the poor. *Forbes Magazine*, 26 Sept 1997.

[26] PMC-Sierra. *Diagram of a 10 Gigabit Core Router Architecture*, April 2002. `http://www.pmc-sierra.com/products/diagrams/CoreRouter_lg.html`.

[27] PMC-Sierra. *Diagram of a Sub-wavelength Optical Cross Connect*, April 2002. `http://www.pmc-sierra.com/products/diagrams/SubWavelengthCrossConnect_l%g.html`.

[28] Private communication. Source requested not to be identified, April 2002.

[29] RHK. North american telecom capex to turn up in 2004. Press release #154, RHK, April 2002.

[30] RHK. Various industry reports and industry news, 2002. IN#114, IR#1101, IR#1102, IR#1079.

[31] US Census. Industry quick report. Technical report, US Department of Commerce, 1997.