Internet Problems and Potential

Henning Schulzrinne
Dept. of Computer Science and Dept. of Electrical Engineering
Columbia University
New York, NY 10027
schulzrinne@cs.columbia.edu

March 18, 1997

Abstract

The Internet has created new modes of communication, but is now poised to also serve as the underlying technology for services currently provided by separate networks. However, as much as it offers additional functionality and possibly cost advantages, the provision of telephony and audio/video distribution services requires changes to the current Internet protocols and the Internet access infrastructure.

1 Introduction

The Internet and local area networks have become a central part of the communications infrastructure in commerce, R&D and education. Many in the technical and educational fields can probably function and continue to work reasonably well without an outside phone connection or PBX for a few hours, but are severely inconvenienced if the internal or external networks are unreachable. Engineers and researchers often receive far more email than phone calls, faxes and postal mail combined. However, the importance is still largely limited to the technical community. As Table 1 shows, the overall volume of email is dwarfed by daily postal deliveries and long-distance phone calls.

Table 1: Communications volume per day (United States)

means of communications	year	mio./day
US Postal Service	1995	580
AT&T US phone calls	1995	200
ATM transactions	1993	20
UPS daily deliveries	1995	12
AOL email	11/1996	7
Federal Express	1995	2

	year	source	min./day
long-distance calls, US	1994	[1]	14
America Online (AOL)	1997	[2]	32
local phone calls, US	1994	[1]	38
total phone line usage	1994	[1]	52
Internet	1995	[3]	57
Television, Sweden	1996	[4]	120
television, USA	1996	[4]	240

So far the Internet has thrived on services where there is no ready competition. As we will discuss in the remainder of the paper, there are a number of reasons to make Internet technology the common platform for a truly integrated services network. Traditionally, service integration has taken place at the trench level (several fibers sharing a single duct), at the physical layer (e.g., WDM or SONET) or at the link layer (ISDN and ATM). Service integration at the network layer offers the same packet-based service platform with different lower-layer technologies spanning about nine orders of magnitude in link speed.

In this paper, we address some of the challenges of turning the Internet into a true "integrated services digital network" or "full service network" that can carry traditional Internet services as well as telecommunication services such as telephony and distribution services such as radio and television. A related paper [5] contains additional data. We will emphasize two related kinds of services, namely continuous-media services delivering continuous streams of multimedia data, such as streaming audio and video, and a subclass of continuous-media, real-time services, which adds end-to-end delay requirements. Both types of services cannot be congestion controlled on the same time scales as, say, bulk data transfer, but their bit rate can be controlled [6, 7, 8].

2 Telephone Services

2.1 Motivation for Internet Telephony

Compared to the current circuit-switched network controlled by a separate signaling network, using Internet technology to provide telephony services has a number of advantages. Audio quality can span the range from communication (2.4 kb/s) to professional quality (256 kb/s and above), while silence suppression reduces bandwidth requirements, particularly in group communications. Conference participants can be identified as they speak. It is far easier to migrate to multimedia services that incorporate shared applications and provide "intelligent-network" facilities. Having access to a data infrastructure simplifies computer-telephony integration. Sharing local and access facilities increases utilization (see below).

It should be noted that the current transoceanic cables already perform compression of audio down to 16 or 32 kb/s, remove silence periods and separate fax traffic, so that the bandwidth savings on international circuits due to Internet telephony will be minor.

The principal disadvantages of Internet telephony are the higher complexity and cost of end systems, the difficulty of powering them during power outages and the lack of reliability. There are a number of possible scenarios for the deployment of Internet telephony, listed here in an estimate of time to widespread use. Clearly, these scenarios can exist side-by-side.

PBX replacement: Institutions and corporations that already operate a high-speed switched network infrastructure and where most employees have a PC or workstation on their desk can readily replace the existing PBX with packet telephony. For a switched infrastructure, voice is only a small fraction of the installed bandwidth, so that resource reservation is not necessary. Even with packetization delays, the end-to-end delay will remain below 40 ms. The gateway between the PSTN and the LAN can be readily implemented in a workstation or PC, particularly since high compression rates are not necessary for LAN applications.

Wide-area transport: In this model, a long-distance provider carries voice calls as packets, but offers a standard POTS interface to the customer. This may be advantageous as a high-speed packet switch with capacity equivalent (about 10 Gb/s) to a standard trunk switch costs significantly less than a circuit switch.

Packet to network termination: In this model, the whole network infrastructure up to the customer network termination is packet-based, but a simple "packet voice" adapter based on a single-chip DSP translates the packet voice stream into a standard two-wire telephony interface. Per-user power requirements are only about 80 mWh/day and 30 W for the switch.

End-to-end Internet telephony: Packetized voice reaches the end system, which may be either a dedicated packet telephone, a network computer or a full-function workstation. Adding packet processing to a microprocessor-equipped phone adds little cost. (Standard ISDN phones cost as much as a low-end Pentium-based computer.)

The efficiency that packet voice gains by compression and voice activity detection (or silence suppression) is partially lost by packetization overhead, as demonstrated in Table 2. Note that the packetization overhead on slow-speed serial lines can be reduced to one or two bytes per packet via header compression [9].

While the transport of real-time data across the Internet is reasonably well understood and standard protocols such as RTP [10] and RSVP [11] have been established, the signaling for setting up calls is still in flux. H.323 [12] [1] provides simple telephony services, but does not support multicast, does not provide advanced services such as personal mobility and, because it is based on ISDN signaling, is rather complex, incurring potential wide-area signaling delays of upward of one second. The Session Initiation Protocol [13] leverages the Internet infrastructure and web functionality.

[1] comprising H.225.0, H.245 and H.323

Table 2: Total bit rate for packet voice, with voice activity suppression (VAD)

protocol stack	units	G.723.1				G.729	
codec bit rate	kb/s	5.6		6.3		8.0	
frame duration	ms	30	60	30	60	10	60
frame size	bytes	20	40	24	48	10	60
ATM CBR (no VAD)	kb/s	6.2	6.2	7.1	7.1	9.0	9.0
ATM VBR, AAL5, RTP	kb/s	10.1	6.5	10.6	7.0	25.5	7.9
IPv4, UDP, RTP	kb/s	8.0	5.3	8.5	5.9	20.0	6.7

3 Multimedia Distribution Services

Given the ubiquity and low cost of distributing radio and television over the air, via cable or satellite, replacement by an IP-based infrastructure seems less imminent, with parallel delivery the more likely scenario. The Mbone experience [14] shows that there is a demand for wide delivery of specialized programming such as shuttle launches, seminars and specialty cultural fare, each with audiences measured in the tens to hundreds. Delivery via IP multicast facilitates receiver-based processing, such as time-shifting, filtering, creation of custom programs, archival and links to other material. Transmitting all 12,140 AM and FM radio stations in the US via the Internet would require about 680 Mb/s at close-to-FM quality, but clearly, because of content and format duplication as well as local advertisement, this would seem to be unnecessary.

Currently, a number of web sites offer audio on demand, delivered via unicast. However, even a T3 connection can only support about 1,500 listeners at barely acceptable audio quality at that rate. There is a need for the development of near-media-on-demand services on the Internet that utilize banks of multicast addresses to deliver staggered versions of popular programming.

Rather than sharing a common network layer, the DAVIC consortium has proposed a "vertical" architecture, where MPEG content delivery takes place over a variety of media, such as ATM, hybrid fiber-coax or even IP. Control and data services may be carried as MPEG 2 transport stream packets [15], using a variety of RPC mechanisms such as DCE and CORBA. In the author's opinion, this application-centric approach is architecturally flawed, as it imposes complexity[2] for each new service and exposes applications to the details of the underlying network. Given that the same end systems are likely to be used for a range of services and that services such as media-on-demand and conferencing are going to be combined within the same application, this imposes a large implementation burden. The initial notion that systems like set-top boxes would not be capable of handling a full Internet protocol stack or web browser has been disproven by, say, WebTV.

The current division between connection-oriented mobile telephony and low-bandwidth packet radio services is severely restricting the kind of services that can be deployed. A single, flexible-bandwidth service with multicast support would avoid the need for separate receivers for data services, voice and multimedia distribution services.

4 Access

The continuous-media services described above impose new requirements on network access. The current residential Internet access architecture, based on modems, is unsatisfactory in several respects:

- The local phone switches were not designed as leased-line switches. While the average local voice call lasts 2 to 5 minutes [16], the average Internet call already lasts for 17 to 21 minutes [16]. Table 1 shows that if broadcast services were to be carried on the Internet, several additional hours could be added to phone line usage.

- The current standard flat-rate fee of about US$20 for unlimited Internet access is based on a statistical sharing model where Internet connection, leased-line, phone company and communication equipment depreciation add to about $6 a month [3]. If subscribers were all on-line during the early evening hours, continuously drawing 56 kb/s of voice or video data, the modem depreciation and ISP business line charges alone would amount to $30

[2] The DSM-CC specification, which does not include the description of MPEG and the RPC protocols, is about 400 pages long.

[3] with the remainder presumably paying for technical support and general business expenses

a month, with an additional $110 for a subscriber's share of the Internet access link. This estimate corresponds roughly with the $75 to $200 per month charged for permanently-connected dial-up business Internet access [4].

- For applications like Internet telephony, remote monitoring and control, the end system should be permanently powered up and connected to the Internet in order not to have to pre-arrange phone calls by email.
- It makes sense to design a network for the predominant traffic type. Data traffic, including fax, has exceeded voice traffic since 1995 [17]. Internet backbone traffic is doubling every three or four months, while long-distance voice traffic is growing by a few percent each year.

5 Quality-of-Service and Pricing

A number of models for providing quality of service suitable for real-time and continuous-media services can be envisioned, enumerated here in rough order of complexity:

Best-effort with traffic engineering: Based on the assumption that turning down "average" requests for service is not really an option for a commercial network that wants to retain customers, the backbone would be dimensioned according to the statistical estimate of customer demand, just like the phone system. This only works if an individual traffic source can only occupy a small fraction of the backbone link bandwidth. Thus, a reasonable approach for customer access speeds beyond dial-up modems might be to allow a subscriber unrestricted TCP access, up to the access bandwidth, but restrict unreserved UDP access to, say, 56 kb/s. This would make reservations and itemized billing unnecessary for most purposes. At least in the United States, telephone access charges are becoming more and more independent of time-of-day and distance [18], approaching flat per-minute fees, with some providers even offering unlimited business service for one day of the week. Unlike phone service, however, it would be easy to simply, say, let the Internet media-on-demand service play, even if there is nobody listening. This could be prevented by volume billing. Already, some T3 services are priced accordingly. This is trivial to bill for, requires no authentication and can be readily verified by the customer.

The case of multicast is more difficult. There is no reason to restrict reception even of high-bandwidth services if a dozen other subscribers on the same access LAN want to receive the same service, since the per-user bandwidth used in the wide area would likely be lower than if the same users were to make use of individualized transmissions.

Commingling TCP and continuous-media traffic within the same queues, whether they operate in tail-drop or random-early drop [19] mode, appears to be sub-optimal, as the deep queues necessary for decent TCP performance lead to long and highly variable queueing delays for continuous-media traffic.

Priority: If continuous-media traffic can be policed so that the sum of the peak rate remains below the link capacity, its delay performance may be greatly improved by giving scheduling priority. This could be easily implemented by having routers inspect the IP type-of-service or precedence field on UDP traffic. Assuming a maximum packet size of 8 kB, a network diameter of 10 hops and OC-48 (2.4 Gb/s) links, the maximum delay variation due to waiting for lower-priority packets is then only 0.3 ms. If continuous-media applications probe for available bandwidth by measuring packet loss, the general loss level will have to rise to statistically observable levels.

Again, charging in such a network is easy and can be based on two simple byte counts.

Class-based scheduling: As a refinement of a simple priority mechanism, various slotted or fair-queueing mechanisms can be employed to ensure each class an appropriate level of service, while protecting it from starvation.

Per-flow reservation: For large flows, i.e., those consuming a significant fraction (say, above 1%) of the link bandwidth, individual reservations are needed to ensure that these flows do not experience excessive loss and delay during their lifetime.

The resource reservation protocol proposed for the Internet (RSVP [11]) offers advantages such as several different reservation styles, receiver-orientation and avoidance of "orphan" reservations, but because of its soft-state mechanism, does not support flows requiring assured reservations. Also, for some applications such as multimedia delivery, sender-based reservations may simplify both the network state and applications designs.

[4] although these charges are probably estimated based on the average usage being less than the modem rate

Currently, RSVP does not support true receiver diversity. However, receiver diversity may be achieved in combination with application-layer entities such as RTP translators that transcode from high-bandwidth to lower-bandwidth media codings.

As an integrated-services network with a large fraction of adjustable-rate data traffic, the Internet has the advantage that the resource estimation for guaranteed services does not have to be very precise, since any left-over capacity can be absorbed by TCP streams.

Early ATM performance papers assumed that video and audio services would require cell loss rates of 10^{-12} to 10^{-9}. The Internet experience has shown that, properly designed, video and audio codecs can deliver very good quality at loss rates of 1 to 10% and higher [20]. This greatly simplifies network design and makes it possible for TCP and continuous-media to be scheduled in the same traffic class.

6 Protocols

The Internet basically operates in four layers, link, network, transport and application-specific protocols, with occasional "shim" layers such as PPP [21], SOCKS [22], SSL/TLS [23] or audio/video framing [24] added as needed.

6.1 Network Layer

Part of the technical success of the Internet protocols can be traced to the datagram service model, which allows more demanding services to be layered on top of this most basic of packet services, yet, since it can be supported by any link-layer technology, can hide the details of the technology from the upper layers.

At the network layer, it will be interesting to see whether and how fast IPv6 can displace IPv4. Certainly, the sense of urgency has largely vanished and the motivation for an individual site is not compelling, since almost all advanced services such as resource reservation, security and mobility can be supported, if less elegantly, by IPv4. For real-time packet audio services, the IPv6 header of 40 bytes, even without security extensions, imposes a significant additional cost.

An additional strength, as compared to ATM [25], is the lack of distinction between user-network and network-network interfaces, with many gradations between minimal end systems and full-fledged core routers.

6.2 Transport Layer

The robustness of TCP has managed to paper over many deficiencies of the underlying Internet infrastructure. However, despite extensions such as large windows, selective acknowledgments or improved congestion control algorithms accumulated over the years, having TCP as the only congestion-controlled transport protocol in the Internet is showing its limitations. First, while the window flow control of TCP is easier to stabilize than rate control since it inherently paces the flow of packets into the network, window flow control becomes inefficient when the average per-connection bandwidth drops below one packet per round-trip time. TCP's only feedback about network congestion is packet loss[5], which leads to unfairness for long-distance connections and throughput oscillations. Research into ATM congestion control [26, 27] seems to indicate that explicit congestion indication, preferably indicated as the bottleneck rate, leads to better performance, albeit at the cost of complexity in the network nodes. Providing congestion indication at the network layer would also allow non-TCP applications such as continuous-media services [8] to adapt to network congestion.

6.3 Application-Specific Protocols

Unlike other network architectures, simple applications can be build quickly using Internet application-specific protocols, without requiring a large supporting infrastructure such as ASN.1, ROSE, DCE or CORBA. Being text-based, they fit nicely with modern scripting languages such as Perl, Tcl, Python or Visual Basic, again accelerating deployment. It is also easier to debug applications and write test tools.

Unfortunately, although there is a large overlap in functionality between the primary Internet data applications such as electronic mail (SMTP, POP and IMAP), network news (NNTP), remote login (rlogin and telnet), file transfer (ftp) and web browsing (HTTP), the application-layer protocols have developed somewhat independently, requiring new efforts to add security, internationalization and feature negotiation for each. With the possible exception of

[5] with exceptions such as TCP Vegas

remote login (rlogin and telnet), the application-layer protocols are basically query-response, that is, remote-procedure calls. One mechanism to unify and simplify the application-specific protocols would be to adopt modern RPC for this, but this would defeat the simplicity and low cost-of-entry advantages. Another is to explore how the most "modern" of these protocols, HTTP, can be extended and generalized to serve as the building block for a range of related protocols, with common facilities to have end points negotiate about extended feature sets, character sets and available resources, authenticate each other and handle charging for application services. Currently, two protocols for initiating and controlling continuous-media sessions, SIP [13] and RTSP [28], are attempting to leverage the best features of HTTP, even though their services have little to do with retrieving web pages.

6.4 Applications

At the application layer, concerns about security dominate. This is less a problem of a lack of protocol support, with usable services now available at the network, transport and application layer, as the lack of infrastructure, e.g., for issuing personal certificates.

As the web is becoming the universal user interface [29] to operating systems and network devices, how to extend trust to applications residing on remote systems becomes more of an issue, as the simple local/remote distinction becomes less useful. Having applications with carefully circumscribed behavior such as Java seems to offer a finer-grained protection than simply having the author sign code, although the latter is useful to guard against "well-known" applications being modified.

On the PC platform, the lack of privilege distinctions exacerbates the problem. With some delay, Windows-based OS will have to acquire the multi-user functionality that workstation OS have had more two decades.

The current mechanism for addressing resources and users in the Internet are somewhat limited. Except for hierarchies of general-purpose caches, which do not work too well for continuous-media services and have billing and accounting problems, we cannot easily obtain the closest copy of a particular resource. On the other hand, addressing email to users not affiliated with a company or university is awkward. Traditional notions of geographic addressing have little meaning in the Internet, and neither provider-based nor vanity domains seem to be satisfying solutions[6].

7 Summary

The Internet protocols have, despite their occasional lack of elegance and functionality, managed to become the lingua franca of inter-organizational data communication. Others promise to replace, at least partially, existing telecommunications services that have not changed much in a century. The Internet as infrastructure is slowly evolving, as more of our communication needs depend on it, into *the* primary channel of electronic communication and commerce. However, as this article has attempted to point out, a range of open problems, requiring research, protocol development and infrastructure enhancements, remain before the Internet and the associated protocols can indeed become the universal communication service.

References

[1] Industry Analysis Division, Common Carrier Bureau, "Trends in telephone service," tech. rep., Federal Communications Commission, Washington, D.C., May 1996.

[2] L. Zuckerman, "America Online moves to placate angry users," *New York Times*, Jan. 1997. CyberTimes.

[3] Emerging Technologies Research Group, "The american internet user survey," survey, Find/SVP, New York, Oct. 1996. Reported in *Tampa Tribune*, January 12, 1996.

[4] D. Lavin, "Tailoring new products is becoming more complex," *Wall Street Journal*, Aug. 1996. European Technology Section.

[5] H. Schulzrinne, "Re-engineering the telephone system," in *Proc. of IEEE Singapore International Conference on Networks (SICON)*, (Singapore), Apr. 1997.

[6] Some solve this problem by having a globally unique identifier.

[6] J.-C. Bolot, T. Turletti, and I. Wakeman, "Scalable feedback control for multicast video distribution in the internet," in *SIGCOMM Symposium on Communications Architectures and Protocols*, (London, England), pp. 58–67, ACM, Aug. 1994.

[7] I. Busse, B. Deffner, and H. Schulzrinne, "Dynamic QoS control of multimedia applications based on RTP," *Computer Communications*, vol. 19, pp. 49–58, Jan. 1996.

[8] R. Bollow, "Video transmission using the available bit rate service," Master's thesis, Berlin University of Technology, Berlin, Germany, Jan. 1997.

[9] S. Casner and V. Jacobson, "Compressing IP/UDP/RTP headers for low-speed serial links," Internet Draft, Internet Engineering Task Force, Nov. 1996. Work in progress.

[10] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: a transport protocol for real-time applications," Request for Comments (Proposed Standard) RFC 1889, Internet Engineering Task Force, Jan. 1996.

[11] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource reservation protocol (RSVP) – version 1 functional specification," Internet Draft, Internet Engineering Task Force, Nov. 1996. Work in progress.

[12] International Telecommunication Union, "Visual telephone systems and equipment for local area networks which provide a non-guaranteed quality of service," Recommendation H.323, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, May 1996.

[13] M. Handley, H. Schulzrinne, and E. Schooler, "SIP: Session initiation protocol," Internet Draft, Internet Engineering Task Force, Dec. 1996. Work in progress.

[14] H. Eriksson, "MBONE: The multicast backbone," *Communications ACM*, vol. 37, pp. 54–60, Aug. 1994.

[15] ISO/IEC, "Information technology – generic coding of moving pictures and associated audio informaiton – part 6: extension for digital storage media and control," Draft International Standard ISO 13818-6, International Organization for Standardization ISO/IEC JTC1/SC29/WG11, Geneva, Switzerland, Nov. 1995.

[16] Bell Atlantic, "Report of Bell Atlantic on internet traffic," FCC report, Bell Atlantic, Mar. 1996.

[17] A. Dorman (Pacific Bell), "Simba newswire," Mar. 1995.

[18] Anonymous, "The death of distance – a survey of telecommunications," *The Economist*, Sept. 1995.

[19] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, pp. 397–413, Aug. 1993.

[20] J.-C. Bolot and A. V. Garcia, "Control mechanisms for packet audio in the internet," in *Proceedings of the Conference on Computer Communications (IEEE Infocom)*, (San Fransisco, California), Mar. 1996.

[21] W. Simpson, "The point-to-point protocol (PPP)," Request for Comments (Standard) STD 51, RFC 1661, Internet Engineering Task Force, July 1994. (Obsoletes RFC1548).

[22] M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, and L. Jones, "SOCKS protocol version 5," Request for Comments (Proposed Standard) RFC 1928, Internet Engineering Task Force, Apr. 1996.

[23] A. Freier, P. Karlton, and P. Kocher, "The TLS protocol," Internet Draft, Internet Engineering Task Force, Dec. 1996. Work in progress.

[24] H. Schulzrinne, "RTP profile for audio and video conferences with minimal control," Request for Comments (Proposed Standard) RFC 1890, Internet Engineering Task Force, Jan. 1996.

[25] J. P. G. Sterbenz, H. G. Schulzrinne, and J. D. Touch, "Report and discussion on the IEEE ComSoc TCGN gigabit networking workshop 1995," *IEEE Network*, vol. 9, pp. 9–21, July/August 1995.

[26] K. W. Fendick, "Evolution of controls for the available bit rate service," *IEEE Communications Magazine*, vol. 34, pp. 35–39, Nov. 1996.

[27] D. Sisalem and H. Schulzrinne, "Congestion control in TCP: Performance of binary congestion notification enhanced TCP compared to Reno and Tahoe TCP," in *International Conference on Network Protocols (ICNP)*, (Columbus, Ohio), pp. 268–275, Oct. 1996.

[28] H. Schulzrinne, A. Rao, and R. Lanphier, "Real time streaming protocol (rtsp)," Internet Draft, Internet Engineering Task Force, Feb. 1997. Work in progress.

[29] H. Schulzrinne, "World-wide web: Whence, whither, what next?," *IEEE Network*, vol. 10, pp. 10–17, March/April 1996.