RE-ENGINEERING THE TELEPHONE SYSTEM

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

Abstract

The evolution of the Internet offers the opportunity to redesign how basic communication services are provided. All services, including (multimedia) telephony, radio and television, can be based on a common, packet-switched, connectionless infrastructure of the Internet. Intelligence is placed into end systems for scalability and controllability. This paper discusses the protocols and technologies that might be used for such a full-service Internet, in particular the role of the access network, resource reservation, and signaling.

1 Introduction

The current Internet and intranets are making the transition from being a convenient additional means of communications that one can easily do without to an essential communication tool. 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.[a] Many engineers and researchers now receive far more email per day, often more than a hundred messages, than phone calls, faxes and postal mail combined. The importance is still largely limited to the technical community, however. As Table 1 shows, the volume of email is dwarfed by daily postal deliveries and long-distance phone calls. Also, only 32% of those who go online say they would miss services[2].

Table 1: Communications volume per day (United States)

means of communications	year	millions/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

[a] Compare this to the effect of a prolonged phone outage documented by Wurtzel and Turner[1].

However, just as telephone and fax have started to displace international postal mail, the Internet could have similar displacement effects as it becomes ubiquitous. (The amount of international mail from the US has been *declining* since 1992 by about 5 to 7% annually, while international call volume rose by 32% between 1992 and 1994.)

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 describe the advantages and implications of using the Internet as a "full-service network". The provision of telephony services is discussed in Section 2, while radio and (cable) television are briefly covered in Section 3. Providing these services accentuates the inadequacy of how residential users access the Internet; Section 4 points out alternatives. A suitable protocol architecture for real-time Internet services is summarized in Section 5. Reliability is crucial for these commodity services, as pointed out in Section 6.

2 Telephony

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:

Compression: Internet telephony allows the parties to use the encoding most appropriate for their quality needs. They may, for example, decide that for an international call, they would trade lower cost for full toll quality, while a reporter calling in her story to the radio station may go for full FM quality with little regard for price. Even without quality degradation, 5.3 kb/s (G.723.1) to 8 kb/s (G.729) are sufficient to support close to toll quality as opposed to 64 kb/s for the current landline phone network. This flexibility also has the advantage that during severe network overload, e.g., after a natural catastrophe, telephone customers can still communicate at about 3 kb/s, increasing network capacity twenty-fold.

Compression benefits the provider if services are price at a flat rate, except that it may allow better use of other services if the access bandwidth is limited.

These services include emulation of a second phone line on a single access line or non-voice services such as simultaneous use of a subscriber's line for a phone call and web browsing.

Silence suppression: Sending audio as packets makes it easy to suppress silence periods, further reducing bandwidth consumption, particularly in a multi-party conference or for voice announcement systems.

Traffic separation: Sending faxes across a circuit-switched network is rather inappropriate, as this is delay-insensitive, but loss-sensitive traffic. Currently, typical fax machines use only 9.6 kb/s of an access line that could support 56 kb/s or 64 kb/s. Thus, fax traffic should be separated from voice traffic as close to the fax machine as possible and converted into either email messages or a TCP connection[3].

User identification: Standard telephone service offers caller id indicating the number (or, occasionally, name) of the caller, but during a bridged multi-party conference, there is no indication of who is talking. The real-time transport protocol (RTP)[4] used for Internet telephony easily supports talker indication in both multicast and bridged configurations and can convey more detailed information if the caller desires.

Computer-telephony integration: Because of the complete separation of data and control paths and the separation of end systems, computer-telephony integration (CTI)[5] is very complex, with specs[6] running to 3,300 pages. All the call handling functionality can be much more easily accomplished once the data and control path pass through intelligent, network-connected end systems. We will describe such functionality in Section 5.2 below.

Shared facilities: Many corporations and universities already have high-speed local area networks. Given its low bit rate, packet voice and low-bit-rate video can be readily supported on a well-designed (switched) LAN, even without explicit quality-of-service support.

Advanced services: From first experiences and protocols, it appears to be far simpler to develop and deploy advanced telephony services in a packet-switched environment than in the PSTN (public switched telephone network)[9,?]. Internet protocols[7] that support standard CLASS (Custom Local Area Signaling Services)[8] features take only a few tens of pages to specify. They can replace both the user-to-network signaling protocols such as Q.931 as well as the network signaling (ISUP, Signaling System 7) and, through cryptography, can be made at least as secure as the existing network.

In the Internet, application-layer intelligence resides in end systems, which are typically replaced much more frequently and have an higher aggregate processing power than typical telephone switches. It is also easier to deploy services one-by-one, rather than waiting for the whole network to be upgraded. (On the other hand, implementing services in switch adjuncts makes them available immediately to all subscribers, regardless of the intelligence of the end system.) Due to implementation diversity, it may also be less likely that software faults in implementations of new features would bring down the whole network. (The Internet can also be used for service creation in the circuit-switched telephone system[9].)

2.2 Cost

One major disadvantage of Internet telephony is the cost of the end systems. It is hard to build packet voice "telephones" requiring no external power that operate over low-grade twisted pair wires several miles long at the $20 price point of a basic analog phone. However, Internet phones are not restricted to being applications on personal computers. They are natural applications for network computers. Also, a voice-only "packet phones" can be built with a single DSP with on-board A/D and D/A conversion and serial interface. A RS-422 or RS-485 interface can operate over a distance of 4,000 feet at 100 kb/s.

2.3 Delay

Compared to circuit-switched telephony, packet voice incurs additional delays beyond the usual speed-of-light propagation delay of roughly 5 ms per 1000 km of fiber. These additional delays are due to packetization, transmission and queueing:

Packetization delays: Packetization intervals are typically chosen between 10 and 50 ms. Shorter packetization intervals increase the header overhead, but make individual packet losses less audible and easier to conceal. Frame-based (typically, low-bitrate) codecs impose a minimum packetization delay of one codec frame, up to 30 ms, and algorithmic look-ahead of up to 7.5 ms.

Transmission delay: Each router or switch hop introduces on the order of 10 μs of transmission delay, assuming a line speed of 155 Mb/s (OC-3) and 200-byte packets.

Queueing delay: The queueing delay depends on the scheduling algorithm chosen. There are at least four choices: best-effort, priority for voice, single guaranteed class for all voice sources and per-source queueing. Best-effort service, the current policy, offers no delay bounds and no loss guarantees, but can work reasonably well in lightly-loaded high-speed networks such as most LANs.

Giving voice and other CBR-like traffic priority ensures that voice traffic does not suffer congestion loss and limits the delay to waiting for any in-progress transmissions. If we do not allow preemption of lower-priority classes and assume a maximum packet size of 8,000 bytes, each router with OC-3 links adds at most 0.4 ms of delay. This estimate is likely to be on the high side, since backbone speeds are likely to be at least OC-12 within a few years and Ethernet LANs limit packets to 1,592 bytes.

An alternative which avoids the problems introduced by simple priority schemes is to treat all voice connections as a single flow for purposes of reservation and scheduling, allocating, by some variation of weighted fair queueing, the appropriate aggregate bandwidth to that class. The delay bound is similar to the one discussed in the paragraph above. Unfortunately, this type of allocation is not supported by the current version of RSVP.

Finally, each flow can be admitted and schedule individually, with the concomitant overhead, but the same delay bound.

Given the per-hop delays, total end-to-end delays for the 40 ms packetization interval typical for low-bit rate voice should consist of the propagation delay plus about 44 ms, of which 4 ms are the total transmission delays for a ten-hop path. For sample-based codecs only, IP-over-ATM or native ATM systems can reduce the packetization delay by transmitting ATM cells as they are filled rather than waiting for a whole packet.

For end-to-end packet telephony, there are no transitions between two-wire and four-wire cabling and thus, no hybrids. This avoids the electrical echo introduced by hybrids that otherwise limits voice circuits to end-to-end delays of about 45 ms without echo cancellation[11][12]. This leaves the acoustical feedback at the handset, which is generally negligible, and allows serviceable connection quality with delays of several hundred milliseconds[13]. Naturally, if speakerphones are used, echo cancellation is required.

2.4 Interoperation with Circuit-Switched Telephony

It is clear that the 700 million or so telephones in the world will not be converted any time soon to packet telephones. (There are only about 200 million computers of all kinds in the world.) Dumb, cheap end systems will continue to be appropriate in many circumstances. Thus, interoperation between "classical" telephony and packet telephony is an important consideration. This can take place in various ways, as

Figure 1: Possible interactions between Internet telephony and POTS

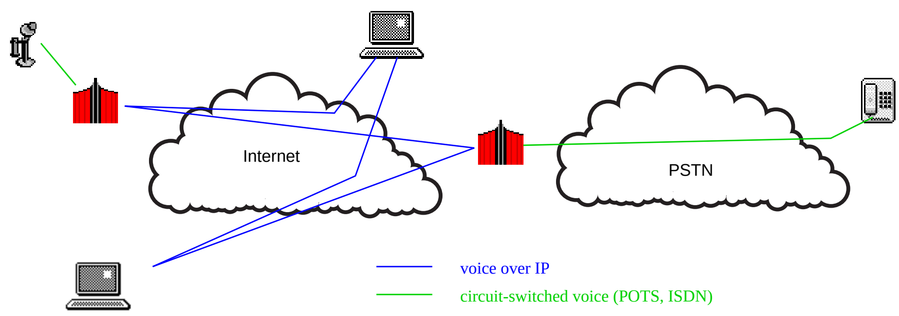

shown in Fig. 1. We can distinguish the following cases, which can be combined as needed:

end-to-end packet: End systems such as network computers, dedicated "Internet phones" or PCs packetize audio; the packets are delivered to one or more similar end systems for playback.

tail-end hop off: Packet networks are used for long-haul voice transmission, while standard circuit-switched voice circuits connect the CPE (telephones) to the packet telephony gateways. This can be used both for individual voice circuits as well as for PBX interconnect. Tail-end hop off allows to bypass long-distance phone companies as well as to connect POTS (plain old telephony service) devices to packet audio end systems.

local packet delivery: Voice is generated by packet audio end systems, but carried as circuit-switched voice over leased or public facilities in the wide area. An example is the "packet PBX" connecting to the PSTN.

term	originating system	wide-area	local-loop termination
	POTS	packet	POTS
	POTS	packet	packet
	packet	packet	packet
tail-end hop off	packet	packet	POTS
	packet	POTS	packet
	packet	POTS	POTS

Table 2: Intensity of communications activities

	year	source	minutes/day
long-distance calls, US	1994	[14]	14
America Online (AOL)	1997	[15]	32
local phone calls, US	1994	[14]	38
total phone line usage	1994	[14]	52
Internet	1995	[16]	57
Television, Sweden	1996	[17]	120
television, USA	1996	[17]	240

3 Radio and Television

The FCC reports that as of December 1996, 12,140 AM/FM radio stations (4,857 AM, 5,419 commercial FM, 1,864 educational FM) are broadcasting in the United States. In other countries, particularly of smaller area and with government-run or public stations, the number of radio stations is much smaller, for example, only about 100 major stations in Germany. If all such stations were to be made available via Internet multicast at FM quality (56 kb/s), this would take 680 Mb/s. Since many of these stations broadcast identical programs, the likely number of channels transmitted nationwide is much lower. More realistically, offering the 31 channels of audio carried by the DirecTV satellite service would add only 1.7 Mb/s of Internet traffic, while the 45 FM channels in New York City would add 2.5 Mb/s.

Carrying radio services over the Internet allows improved directory services, easy addition of side information such as content labeling, bandwidth diversity for different kinds of programming as well as the reception of more diverse programming particularly in more sparsely populated areas. Labeling also makes it easy to have receivers assemble custom programs for listening at some later time ("time-shifting"). (Instead of constructing a special-purpose digital radio system it would seem to be preferable to have a general packet radio, with some fraction of the bandwidth set aside for distribution services, both audio and other content types.)

Similarly, given enough bandwidth, there are substantial advantages of distributing entertainment video, both broadcast and video-on-demand services, using Internet protocols, in particular MPEG over RTP[18]. This, in combination with an Internet-based stream control protocol such as RTSP[19], appears to be both more flexible and simpler than using MPEG as a new network protocol and DSM-CC as a control protocol.

4 Internet Access

The services just described 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 [20] or 4.28 minutes according to [21], the average Internet call already lasts for 17 to 21 minutes [20], depending on the phone company surveyed. This has already lead to isolated inability to obtain dial tone in some central offices. As shown in Table 2, if carried by the Internet, telephony and broadcast services can add several hours to phone line usage.

- The current standard flat-rate fee of US$19.95 for unlimited Internet access is based on a multiplexing model [22,23] where 200–300 concurrent users share a T1 access line (costing about $3000 per month), for an average rate of 5 kb/s to 7 kb/s. A modem is provisioned for every 10–15 customers, depreciated at about $10/month, with ISP line charges of about $20. Again, with continuous use of the Internet for video entertainment, interactive games or telephony increases, Table 2 suggests that during the busy hours of the early evening, almost every subscriber will be trying to connect to the modem pool. Also, the average data rate will exceed the per-user rate.

  For example, a January 1997 study commissioned by AT&T Worldnet reports that between 6 pm and 9 pm, users manage to connect to a modem 93.4% of the time. Users calling AOL during the same hours, meanwhile, manage to connect just 36.7 percent of the time [24].

- 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.

- Since data traffic, including fax, has exceeded voice traffic since 1995 [25], it makes sense to design the network for the predominant traffic.

It was hoped that ISDN might improve data access, as it avoids some of the signaling difficulties of analog POTS and offers higher speeds, yet configuration complexity, hardware costs, the necessity to maintain an analog line for communications during power outages, comparatively high tariffs, as well as the decreasing gap in speed between analog modems and a single ISDN basic rate channel have limited the appeal of ISDN, at least in the US. Clearly, ISDN does not address the congestion problem at the local office.

Table 3: Cost comparison of switches and routers

switching method	ports	capacity (Gb/s)	cents/kb/s	interface
10BaseT Ethernet hub		0.12	0.5	
10BaseT Ethernet switch		0.24	2.0	
100BaseTX Ethernet switch		0.80	1.0	
router		2.1	16.0	
local ATM switch		2.48	1.6	
PBX (256 lines)		0.02	218.	140
Lucent 5ESS local (no AIN)	5,000	0.32	469.	300
Lucent 5ESS local (AIN)	20,000	1.28	273.	175
Lucent 4ESS toll (100k lines)	6.40	7.8		

From an economical standpoint, common data switches are much more cost-effective for switching bits than either PBXs or traditional telephone switches. A rough estimate of the per-kb/s cost of switching is given in Table 3. Note that costs for the 5ESS and other switches are hard to quantify since the vendors do not generally release price lists. [b].

Figure 2: Residential Internet access

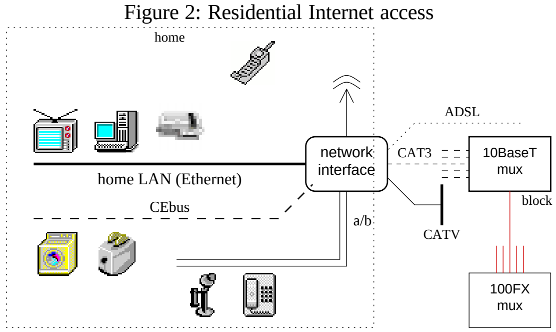

Given these considerations, traffic should be packetized as close to the end user as possible. Thus, a suitable architecture for residential access to the Internet should look similar to the current corporate LAN environment, as sketched in Fig. 2. The figure indicates three alternatives for access, namely ADSL, CATV[26], and Ethernet. The first two have been explored extensively elsewhere. However, for densely populated

[b] e.g., 5 5ESS switches with 100,000 lines for \$20 mio quoted in a February 1991 AT&T press release, 6,000 lines for \$7 mio, but including a satellite ground station, in a May 1991 press release. The PBX costs are based on a mid-size Siemens PBX purchased in 1996.

areas or apartment buildings, an Ethernet-based approach may be more cost-effective. Since the median loop length between network termination and central office of 1.7 miles (9 kft) (1973)[12] exceeds twisted-pair Ethernet specifications, a direct connection between homes and the central office is not possible. Instead a more layered approach needs to be taken. For example, several dozen homes or apartments would be connected to an Ethernet switch or hub, located no more than the CAT-3 cabling distance limit of 328 feet from the network termination unit. The switches would, in turn, connect through fiber to the neighborhood switch. This architecture has the advantage that a mix of low-bandwidth and high-bandwidth customers can be accommodated without running additional wires. Since switch costs are dominated by interface counts rather than bandwidth, this mechanism offers much higher per-user bandwidth (particularly peak bandwidth), yet switching costs are similar to today's.

Ethernet appears advantageous as the local access technology since PC interfaces are cheap, operate over a variety of media and, unlike ATM, they allow easy addition of more devices on a multiple-access LAN.

Even with packet connectivity to the individual residence, easy connection of existing telephone equipment is imperative. Thus, Fig. 2 shows the network termination unit with a built-in a/b (two-wire) telephone interface. This could be readily implemented within a single DSP that would act as a simple packet voice module. It would also implement DTMF recognition for user-to-network signaling based on, say, SIP[7], described briefly below.

5 Protocol Architecture

We now present an overall Internet protocol architecture that can support telephony and other continuous-media services such as "Internet radio" and "Internet TV"[27].

5.1 Data Transport

For transporting real-time data across the Internet, the accepted end-to-end protocol is RTP[4,28]. It is also used by the ITU-T H.323 teleconferencing recommendation. RTP is a thin protocol providing support for applications with real-time properties, including timing reconstruction, loss detection, security and content identification. RTCP provides support for real-time conferencing for large groups within an internet, including source identification and support for gateways (like audio and video bridges) and multicast-to-unicast translators. It offers quality-of-service feedback from receivers to the multicast group as well as support for the synchronization of different media streams.

While UDP/IP is its initial target networking environment, efforts have been made to make RTP transport-independent so that it could be used, say, over CLNP, IPX or other protocols. RTP is currently also in experimental use directly over AAL5 using native ATM services. This protocol stack is shown in Fig. 3.

Figure 3: RTP-based protocol stack

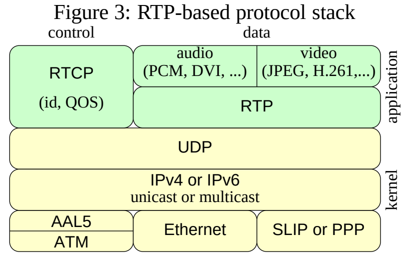

For low-speed links and highly compressed audio, the combined stack consisting of IP, UDP and RTP add 40 bytes to every packet, while 20 ms of 8 kb/s audio only take up 20 bytes. Thus, just like for TCP, header compression is desirable[29].

As an end-to-end protocol, RTP cannot guarantee a certain quality of service. The resource reservation protocol, RSVP [30], might be used to allocate resources to either individual streams or a group of streams. Unlike in the telephone network, establishing connectivity and allocating resources are two distinct operations. Typically, the caller would first "ring" the callee. Only if the called party wishes to communicate and once a set of media has been agreed upon, would RSVP be used to reserve resources. This runs the risk of being denied the necessary network resources, but in any network dimensioned for telephone service, voice calls should be blocked extremely infrequently. Given that the resource needs are only known once the parties agree on the media and their quality, this order is also the only feasible one.

Since packet audio flows have a relatively low, constant bit rate, with deterministic arrival patterns, it may make more sense to treat all voice calls on a link as a single stream for scheduling rather than managing several thousand individual reservations, with the attendant state space and refresh overhead. As indicated earlier, it may be sufficient to simply give this class of service priority. In some cases it may also be advantageous for the end systems to reserve only the minimum necessary bandwidth, and obtain additional throughput and improved audio quality through best-effort flows.

5.2 Signaling

In the architecture described here, the traditional telephony signaling protocols, Q.931 for ISDN user-to-network signaling and SS7 for network-to-network signaling, are replaced by a single, much simpler signaling protocol. One such solution, the session initiation protocol (SIP) [7], can establish multimedia conversations with one or more parties. Instead of telephone numbers, it uses addresses of the form user@domain or user@host. In many cases, this address would be identical to a person's email address. Table 4 compares the properties of email addresses to standard telephone numbers.

Table 4: Comparison of (U.S.) telephone numbers and email addresses

feature	phone number	email
mnemonic	no (except for some 1-800)	name@organization
multiple per person	no	easy
avg. characters	≈ 12	22
location-independent	1-700	yes: j.doe@ieee.org
carrier $\neq$ naming	maybe	yes
directory	411, 1-555, switchboard.com	LDAP[31]

(Note: average email address size is estimated from the IEEE ComSoc TCCC mailing list, an international list with about 500 members.)

SIP offers the standard PBX or CLASS functionality, including call forwarding, call waiting, caller ID, call transfer, camp-on[c], call park[d], and call pickup[e]. Many of these features actually require no signaling support at all, but can be implemented by end system software. SIP is designed as a variant of HTTP/1.1[32], which allows easy reuse of HTTP security and authentication, content labeling and payment negotiation features.

We used the SIP features to implement a calendar-based call handler. The call processing software accesses a user's personal appointment calendar and answers the phone accordingly. The user can define categories of callers and preset, based on the calendar entry, whether and where their calls are forwarded. The information released to the caller if calls are not forwarded may range, for example, from "is currently not available" to "John Smith is in a meeting until 3 pm in room 5621 with Jane Doe", depending on the caller's identity. In the near future, this will be integrated with the call processing language, a state-based scripting language that allows to construct voice-mail systems or automatic call handling systems in a few lines of code. It also manages the translation between ISDN calls and Internet telephony calls.

[c] "Camp-on allows an attendant-originated or extended call to a busy single-line voice station to automatically wait at the called station until it becomes free while the attendant is free to handle other calls." ?

[d] "Allows user to put a call on hold and then retrieve the call from another station within the system". ?

[e] "Allows stations to answer calls to other extension numbers within the user specified call pickup group" ?

6 Reliability

Clearly, if Internet technologies are to play a larger role in providing telephony service, they have to offer at least the same reliability and manageability as the older circuit-switched technologies they are to replace. The target is set fairly high: a typical local telephone switch is out of service for about 120 seconds a year, of which 25% are for outages less than 2 minutes. On average, a subscriber can expect his or her line to be unavailable for 85 seconds during the year, of which 34 seconds are scheduled outages, where only outages of more than two minutes are counted.[f]

Unfortunately, only limited reliability information is available for online and Internet services. America Online reports[33] scheduled and unscheduled outages of 1% or 88 *hours* a year for 1996, down from 3.5% or 307 hours the year before.

A more packet-oriented architecture, as presented earlier, would remove this restriction, as well as modems that are sometimes subject to locking up or not releasing a line. Typical local network switching equipment has actualized MTBFs of 170,000 to 210,000 hours (19 to 24 years); from anecdotal evidence, hubs and Ethernet switches rarely fail as a whole, rather, individual interfaces may.

Recent reported large-scale Internet outages were due to either misconfiguration, as in the case of the AOL BGP router collapse in 1996, or a local power failure without adequate uninterruptable power supplies, as when a major POP on the Stanford University campus was out of service for a better part of a day. Many of the Mbone routing failures are due to router misconfiguration, for example injecting all unicast routes into the multicast routing protocol. Also, it is clearly harder to maintain telephone-level uptime when traffic doubles every few months and host counts double yearly.

However, there are a number of obvious improvements that are necessary for Internet services to approach telephone-level reliability:

- Software upgrades should be possible without taking down a router or switch.
- Router configuration must be made simpler, with checking against local rules that make catastrophic failure less likely.
- Closer integration of network functionality into the operating system should reduce end-system difficulties. Much of the complexity of configuring current PCs for Internet usage appears to stem from having to configure the modem and multiple protocol stacks. Widespread use of DHCP[34] and IPv6 autoconfiguration[35], as well as eliminating the modem, should make the network invisible.

[f] These numbers are drawn from the FCC quality of service reports for 3Q93 through 3Q95 and are for USWest, one of the regional Bell operating companies (RBOCs) with about 13.8 million access lines. It appears that the difference between the per-switch and per-line figures largely explained by the two-minute threshold. An average switch serves 8,600 lines.

- While backbone networks feature redundant links, POPs and access links are often single point of failures. Also, different providers often peer at only a small number of points. Particularly the latter problem must be remedied to increase the number of true end-to-end alternate paths.

In the long run, tools like traceroute and ping, as well as relatively simple management protocols (SNMP) and applications with built-in reporting mechanisms such as provided by RTCP probably make Internet service more manageable than many traditional POTS installations.

A particular problem with packet telephony is the need for power at the end systems, and, if the architecture of Fig. 2 is adopted, at the various multiplexing points within the network. Fortunately, end systems can be build so that they consume almost no power in standby mode and should be able to function on a small rechargeable battery for days, given the current operating times of small cellular phones. Similarly, a typical 24-port Ethernet switch consumes about 30 W of power, so that it can be operated with a typical 1 kWh sealed lead-acid battery for more than 30 hours.

7 Summary

Internet services have the promise of being the foundation of an integrated services, packet-switched networks that delivers not only web pages and email to its users, but also replaces parts of the telephone system or cable television. Within the last few years, many of the necessary protocols and architectures have emerged to realize this vision. However, the most important factor may not be protocols, but new residential access methods, increased reliability and sufficient backbone capacity.

1. A. H. Wurtzel and C. Turner, “Latent functions of the telephone: what missing the extension means,” in *The Social Impact of the Telephone* (I. de Sola Pool, ed.), pp. 246–261, Cambridge, Massachusetts: MIT Press, 1977.
2. Times Mirror Center for The People and The Press, “Survey about computer use,” Oct. 1995.
3. C. Malamud and M. Rose, “Principles of operation for the TPC.INT subdomain: Remote printing – technical procedures,” Request for Comments (Experimental) RFC 1528, Internet Engineering Task Force, Oct. 1993. (Obsoletes RFC1486).
4. 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.
5. C. Strathmeyer, “Feature topic: Computer telephony,” *IEEE Communications Magazine*, vol. 34, Apr. 1996.

6. Versit Consortium, “Computer telephony integration (CTI) encyclopedia,” Tech. Rep. Release 1.0, Versit Consortium, Oct. 1996.
7. M. Handley, H. Schulzrinne, and E. Schooler, “SIP: Session initiation protocol,” Internet Draft, Internet Engineering Task Force, Dec. 1996. Work in progress.
8. P. Moulton and J. Moulton, “Telecommunications technical fundamentals,” technical handout, The Moulton Company, Columbia, Maryland, 1996. see http://www.moultonco.com/semnotes/telecomm/teladd.htm.
9. C. Low, “The internet telephony red herring,” in *Proceedings of Global Internet*, (London, England), pp. 72–80, Nov. 1996.
10. C. Partridge, *Gigabit networking*. Reading, Massachusetts: Addison-Wesley, 1993.
11. J. Bellamy, *Digital Telephony*. New York: John Wiley & Sons, 1991.
12. R. F. Rey, ed., *Engineering and Operations in the Bell System*. Murray Hill, New Jersey: AT&T Bell Laboratories, 1984.
13. International Telecommunication Union (ITU), “G.114: transmission systems and media, general recommendation on the transmission quality for an entire international telephone connection; one-way transmission time,” Recommendation G.114, Telecommunication Standardization Sector of ITU, Geneva, Switzerland, Mar. 1993.
14. Industry Analysis Division, Common Carrier Bureau, “Trends in telephone service,” tech. rep., Federal Communications Commission, Washington, D.C., May 1996.
15. L. Zuckerman, “America Online moves to placate angry users,” *New York Times*, Jan. 1997. CyberTimes.
16. Emerging Technologies Research Group, “The american internet user survey,” survey, Find/SVP, New York, Oct. 1996. Reported in *Tampa Tribune*, January 12, 1996.
17. D. Lavin, “Tailoring new products is becoming more complex,” *Wall Street Journal*, Aug. 1996. European Technology Section.
18. D. Hoffman, G. Fernando, and V. Goyal, “RTP payload format for MPEG1/MPEG2 video,” Request for Comments (Proposed Standard) RFC 2038, Internet Engineering Task Force, Oct. 1996.
19. H. Schulzrinne, “A real-time stream control protocol (RTSP’),” Internet Draft, Internet Engineering Task Force, Dec. 1996. Work in progress.
20. Bell Atlantic, “Report of Bell Atlantic on internet traffic,” FCC report, Bell Atlantic, Mar. 1996.
21. M. Mayer, “The telephone and the uses of time,” in *The Social Impact of the Telephone* (I. de Sola Pool, ed.), pp. 225–245, Cambridge, Massachusetts: MIT Press, 1977.

22. D. H. Dennis, “The inet-access frequently asked questions list,” Dec. 1996.
23. R. Stiennon, “So you want to be an internet provider?,” *TechKnow Times*, vol. 2, May 1996.
24. D. Lazarus, “AOL refugees call on ATT WorldNet,” *Wired News*, Jan. 1997.
25. A. Dorman (Pacific Bell), “Simba newswire,” Mar. 1995.
26. C. Bisdikian, K. Maruyama, D. I. Seidman, and D. N. Serpanos, “Cable access beyond the hype: on residential broadband data services over HFC networks,” *IEEE Communications Magazine*, vol. 34, pp. 128–135, Nov. 1996.
27. H. Schulzrinne, “Internet services: from electronic mail to real-time multimedia,” in *Proc. of KIVS (Kommunikation in Verteilten Systemen)* (K. Franke, U. Hübner, and W. Kalfa, eds.), Informatik aktuell, (Chemnitz, Germany), pp. 21–34, Gesellschaft für Informatik, Springer Verlag, Feb. 1995.
28. 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.
29. 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.
30. L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, “RSVP: a new resource ReSerVation protocol,” *IEEE Network*, vol. 7, pp. 8–18, Sept. 1993.
31. W. Yeong, T. Howes, and S. Kille, “Lightweight directory access protocol,” Request for Comments (Draft Standard) RFC 1777, Internet Engineering Task Force, Mar. 1995. (Obsoletes RFC1487).
32. R. Fielding, J. Gettys, J. Mogul, H. Frystyk, and T. Berners-Lee, “Hypertext transfer protocol – HTTP/1.1,” Request for Comments (Proposed Standard) RFC 2068, Internet Engineering Task Force, Jan. 1997.
33. D. Lazarus, “AOL’s email system crashes twice in one day,” *Wired News*, Jan. 1997.
34. R. Droms, “Dynamic host configuration protocol,” Request for Comments (Proposed Standard) RFC 1541, Internet Engineering Task Force, Oct. 1993. (Obsoletes RFC1531).
35. S. Thomson and T. Narten, “IPv6 stateless address autoconfiguration,” Request for Comments (Proposed Standard) RFC 1971, Internet Engineering Task Force, Aug. 1996.