A Taxonomy of Internet Telephony Applications

David D. Clark
Internet Telephony Consortium
Massachusetts Institute of Technology
ddc@lcs.mit.edu

Abstract

There are a broad range of objectives and opportunities that can be lumped under the heading of Internet Telephony, or ITel. This paper identifies a number of criteria that can be used to separate these different ITel applications into classes, including the degree of interoperation with the existing telephone system, and the extent to which the existing Internet must be augmented to support them. Using this framework, the paper concludes that different ITel applications have very different motivations, and have very different implications for industry structure, economics and regulation. The immediate opportunities for ITel involve cost reductions relative to current telephone pricing. The long term trajectory for ITel is to become a new mode of computer-mediated human communication, which will have profound consequences for the telephone industry. This long-term form of ITel will not necessarily grow directly from the products that are being deployed now, but will come from a number of intermediate developments that can be anticipated over the next several years.

Introduction

Internet Telephony, or ITel, can mean a number of different things. It can mean the use of Internet technology to replace a long distance or international provider of traditional telephone service, or an enhanced form of human to human communication based on the computer as the user interface, rather than the telephone. This paper describes the rather wide range of applications that have been called Internet Telephony, or ITel, and tries to articulate the important differences among them, with the goal of organizing them into major classes.

The purpose of organizing ITel applications into classes is that it provides a framework around which to speculate on the broader implications of ITel. The different classes of ITel have very different justifications, and very different implications for the relevant industrial sectors involved, as well as policy-makers and users.

Some ITel applications are focused on a short term cost savings strategy, which may not have strong long-term market viability. However, a possible long-term outcome of the ITel evolution is that people use computers rather than telephones to communicate. This outcome, were it to happen, could trigger a major restructuring of the telephone industry, in which separate firms provide the low-level physical connectivity and Internet service, and the higher-level telephone service itself.

The final speculative form of ITel described above is not practical today, because the necessary supporting features in the Internet are not in place. It is our hypothesis that Internet telephony will evolve as a series of incremental steps. Early variants of ITel will be identified that can be deployed without first requiring as much enhancement of the Internet. These offerings will serve as experiments to prove the market, evaluate demand, explore the desirability of features, and motivate the fuller deployment of enhanced Internet service.

The paper looks at specific examples of ITel applications in order to develop an approach to organizing and classifying them. We begin by proposing a first basis for classification, which is how much interoperation is required between the Internet and the existing telephone system, or PSTN. We then describe in greater detail what is implied by the long term vision of ITel alluded to above. By listing the critical features of this long-term outcome, we develop a further checklist of features and requirements by which we can organize the different ITel applications. We will show that there are ITel applications with different mixes of these features and requirements, which in turn suggests a possible evolutionary path for the future.

How much PSTN? The most important question.

The most significant distinction between the various Internet Telephony application is the question of how much PSTN and how much computer-based telephony is in the scheme. In this paper we identify three important classes of ITel applications.

- Class 1: proposals with the goal of using Internet to provide POTS telephony between existing telephone end-user equipment. Applications of this class require technology for interconnection between the PSTN and Internet networks, but do not require access to computer-based end-nodes, and can often operate across dedicated regions of the Internet.
- Class 2: proposals that require interoperation between the existing telephone and Internet networks, and provide communication between users with either computers or existing

telephone sets as end nodes. This class requires both the interoperation between Internet and PSTN, as well as the use of computer-based end-nodes.

- Class 3: proposals that use Internet-attached computers to provide some form of human communication across the packet-switched Internet. This class, the pure form of Internet-based communication, does not involve any aspect of PSTN interaction, or interworking with telephone end-nodes. This class is 100% Internet, and 0% PSTN.

These three classes are illustrated in figure 1. We will argue that this classification provides the most powerful way of articulating the different broad categories of Internet Telephony.

To illustrate some of the issues, we will first consider an example of a class 1 service, before looking at some of the other ways that applications can be distinguished.

Example 1: Class 1 -- international/long distance POTS using Internet technology

This variant of ITel uses Internet technology to connect into the existing telephone infrastructure as a international or long distance carrier. Customers continue to use their local phone system and telephones, and see this just as a long-distance or international alternative. Since the customers continue to use their existing handsets, the service is still essentially POTS. While some variation in voice quality is possible, the motivation is to deliver a lower-cost variant on traditional POTS telephony, by using Internet technology.

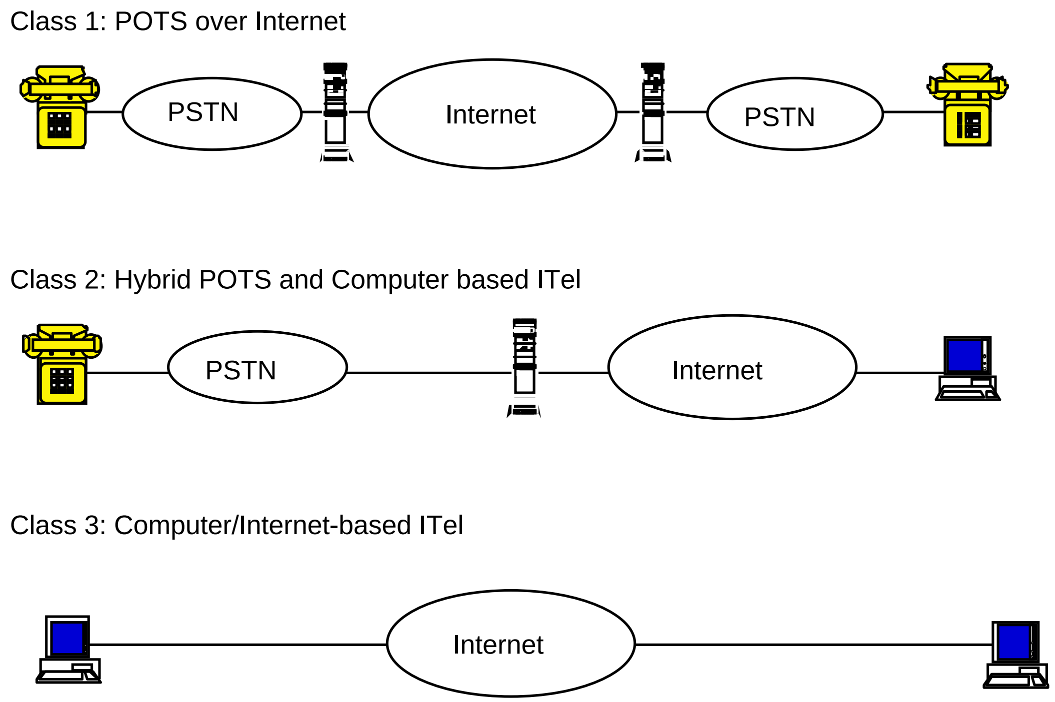

Figure 1: The three classes of Internet telephony.

Reducing the cost of POTS telephony

There are three ways that a lower cost could be realized. First, some of the costs in existing telephony are artificially high, and ITel may be able to side-step these artificial costs. Second, by efficient compression of the voice, costs could reduced at the expense of a somewhat lower sound quality. Third, the Internet technology could deliver a lower intrinsic cost for the same service. We claim that the first factor is actually the only important one

While there has been considerable debate about the relative cost of carrying voice over the Internet and the PSTN, to first order there seems to be no intrinsic cost advantage to carrying POTS-style calls of the same audio quality over the two networks. The argument is that the same circuit is used in both cases, the same compression scheme could be used in both cases if the cost reduction warrants, the compression (if done) is implemented in a component of the same complexity in both cases, and both schemes can take advantage of the statistical nature of talk spurts if warranted. In the long run. the use of Internet (or packet switching in general) does not appear to lead to greatly reduced per-minute costs for carrying a call.[1]

While there may be few long-term intrinsic cost advantages to using Internet technology as a component of the POTS infrastructure, the current telephone prices, with regulated rates and high prices sustained by monopoly players in certain countries, appear to provide a number of options for new entrants to offer a much lower cost alternative to incumbent providers by structuring themselves as Internet providers. In the international market, these options for arbitrage are substantial. In the long distance market in the U.S., prices have already been driven down by competition, so there is less advantage for arbitrage than in the international case. However, the access fees paid by traditional long distance providers to the LECs currently do not apply to the Internet, so long distance provided over Internet avoids theses fees. The motivation for some providers to propose long distance POTS telephone service over the Internet may be the indirect one of forcing the FCC to move on the resolution of the current consideration of local access charges.

Example 2: Class 3 -- long-term Internet based communication

In contrast to the application described above, which is a short-term proposition to exploit price distortions, there is a long term vision of what Internet telephony might be. The speculative end-point of ITel is a general set of applications for computer-mediated human communications. The distinguishing characteristic of this application is not lower cost, but enhanced functions. The computer and the Internet are central to this objective.

The Internet is a natural network for this application (more so than the existing phone system) for several reasons. Packet switching allows these applications to be mixed with others over a common network. Voice can be combined with other modes of communication -- text, video, shared workspace and so on. For telephone-like applications, the Internet can deliver the signaling information all the way to the end-node, so the telephone features can be implemented at the edge of the net. It will permit several calls to co-exist over one physical copper pair (or other

---

[1]In the short run, the Internet solution might have lower costs, because the telephone provider must deal with the depreciation of capital equipment such as circuit switches that makes less efficient use of the same circuit. Additionally, an important second-order effect is that because an Internet infrastructure can support the signaling and operations requirements as well as carry the voice traffic, there may be some significant efficiencies in building a single Internet to support all aspects of the telephone application.

medium), and it supports advanced features such as multicast that permit many-to-many communications.

The computer provides the end-node functionality and the more sophisticated user interface. The computer can be used as a "call manager", keeping track of numbers and unanswered calls, assigning priority to incoming calls and redirecting them as appropriate, and logging and archiving calls. It can personalize the communications service for each user that shares the system, for example providing a different response for business callers, friends, and strangers calling the same location. The computer can assist in lowering costs for communications, by obtaining network service from the lowest cost provider at each instant.

Differentiating characteristics

This form of ITel differs from the class 1 PSTN substitution described above in a number of important ways. These differences will turn out to be important characteristics that can be used to classify variants of ITel generally.

1) Divergence away from POTS functionality

Any system that includes traditional telephone handsets must interoperate in a way that is consistent with the very limited nature of that device. The richness and flexibility of the user interface envisioned by this long-term form of ITel depends on having no (or very little) need for backwards compatibility with the POTS-style service. Only when we reach the class 3 variations (0% PSTN) does the application designer have the option of diverging in a serious sense from the POTS-style interface.

2) Migration of function towards end-node.

A characteristic that is closely intertwined with the above is the extent to which the location at which functions are implemented has migrated from the center of the network to the edge. The traditional telephone system has very primitive end-notes (the telephone) and much intelligence inside the network. The Internet in general represents a different balance, with intelligent end-nodes (the computer) and a simple set of functions inside the switches of the network. This long-term class 3 variant of ITel, to a large extent, will be an application implemented in the end nodes rather than inside the network.[2] Features such as call-waiting and caller-id, which today are

---

[2]The counterbalance to this is that for some Internet applications which have dynamic information shared among several parties, such as multi-player games and MUDs, there may be a pressure to move aspects of the application back into the core of the network. Thus, there are pressures that move function in both directions within the Internet.

implemented inside the telephone network, will just be implemented as software on the user's computer.

3) Ease of use

PSTN telephony is getting more difficult to use, with longer strings of numbers necessary to complete calls, prefix sequences to select long distance carriers (a concept that many consumers find very confusing), and advanced features such as call waiting or voice mail implemented using the telephone keypad as the user interface. The addition of more features can only make this worse. On the other hand, the use of a computer as an alternative user interface might improve the situation. Additionally, if the computer were programmed to automate certain steps (such as selecting the long distance carrier based on the number being called and the current costs from different providers), increased ease of use might be combined with increased value. Ease of use is thus a factor that may get either better or worse, depending on the details of specific offerings.

4) Augments needed to the Internet

In order to bring this Internet based telephony application into existence, it will be necessary for the Internet to evolve in a number of ways: new protocols and technology will need to be developed and deployed. There are five significant enhancements to the Internet that we have identified in order to support Internet telephony:

a) QoS

The term Quality of Service (QoS) is used in Internet design to describe the ability to assign specific treatment within the network to certain flows of packets. For example, a specified flow of packets might receive an assured minimum bandwidth, or a bound on the delay of the delivery. There is work underway now in the Internet engineering and standards community to add QoS support to the Internet protocols.[3] [4]

b) Pricing for enhanced services

---

[3] Note that the term QoS, as used by Internet designers, has nothing to do with what change the user of an application may perceive as a result of using the QoS features of the network. The term QoS describes network-level mechanisms, not user-perceived quality. As a specific example, if a voice application needs a bounded delay on the packet delivery time, this can be achieved either by adding QoS mechanisms to the Internet, or by restricting the offered load (e.g. by blocking voice calls) so that the Internet infrastructure is somewhat underloaded.

[4]Internet services to support real time delivery of audio and similar material are described in [Wroclawski] and [Braden].

If the network infrastructure is enhanced with QoS mechanisms to provide different sorts of service, some controls will be required on the selection of these services by the users. In a private network (a corporate intranet, for example) administrative controls may be sufficient to limit the use of enhanced QoS. However, in the public Internet, pricing seems the obvious mechanism to control the selection of enhanced QoS, since it both limits user consumption and provides increased income to compensate the provider of the service. So pricing mechanisms of some sort seem necessary as a complement to the basic QoS mechanisms. [5]

c) Reliability

Telephony is traditionally associated with a level of reliability greater than that of the current Internet. Providers of Internet technology, e.g. routers, may find themselves pressed to meet these greater expectations. Protocols and mechanisms for requirements such as fault detection and recovery may have to be upgraded. Providers of network infrastructure that supports Internet, such as HFC cable facilities and wireless, may be similarly pressed to improve reliability and availability.

d) "Always on" connectivity

Today, most residential Internet customers are not connected to the Internet at all times, but only when they explicitly dial up. To support general calling patterns among users, for example receiving a call without prearrangement, users would have to be connected to the Internet constantly, so they could receive a call whenever one is placed to them. This pattern of "always on" connection would add to the cost of providing modem-based Internet access, since it would require a modem at the central site be provided for each user, and the modem banks provided by Internet Access Providers represent a significant part of their overall cost. Continuous connection would also increase the load on the telephone switches and trunks of the LECs that provide the dialup circuits.[6] The access technology of the cable industry makes it easy to provide this

[5]A number of approaches to pricing of services in the Internet are discussed in [McKnight &Baily] .

[6]The alternative would be for the Internet Access Provider to dial the customer when access is appropriate. The impediment to this is tariffs. Consumers can purchase flat rate local calling service, but IAPs are requires to purchase business service, which generally has only metered service. So calls *to* the IAP incur no incremental payments to the telephone company, while calls *from* the IAP incur usage charges.

"always on" service over cable, and that industry is not blind to the advantage that this offers them, both in providing ITel services, as well as other applications. [7]

e) Ubiquitous deployment

The term "network externality" is used to describe the attribute that a network is increasing valuable to any one user as more other users are attached. If only a few percent of the population is actually on line, the appeal of enhanced service ITel will be minimal. General purpose ITel depends for its appeal on a sufficiently widespread deployment of the service, the Internet itself, and suitable end-node equipment such as computers to the population.

These five augments to the Internet are all required to support the general form of Internet telephony described in this example. As table 1 illustrates, the short-term application of long distance POTS and the long term application of next generation computer based human communication differ in almost every one of these key attributes.

	Application	Beyond POTS?	End-node fcn?	Ease of use	Internet Augments				
					QoS	Pricing	Reliability	Always on	Ubiquitou
1	LD POTS	no	none	same	no	no	local	no	
3	gen'l ITel	yes	major	better?	yes	yes	yes	yes	

Table 1: Comparison of Class 1 Long Distance POTS and
Class 3 long-term Internet Telephony

In particular, the class 1 variant requires fewer augments to the Internet technology, because one can avoid the need for explicit QoS mechanism in the Internet infrastructure by building a dedicated Internet used only for the long-haul telephony, and controlling the load by blocking calls. By using dedicated capacity, one also avoids the need for billing mechanisms at the Internet level. Finally, since the PSTN is being used to distribute the voice traffic to the end nodes, there is no concern with widespread deployment of Internet, or with providing the user with an "always on" service model.

There are concerns about reliability. Within the specific part of the Internet being used for carriage of the voice traffic, this application does require robust equipment capable of "telephone-grade" reliable operation. However, this requirement only applies to the dedicated

[7] For an alternative discussion of important enhancements to the Internet, including the "always on" feature and several other related requirements, see chapter 2 of [CSTB 96].

region of the Internet being used for the voice. Because few enhancements are needed to existing Internet to implement this application, it is feasible to undertake class 1 ITel today, and there are commercial, Internet-based international providers of telephone preparing to offer service today.

Example 3: Trying class 3 Internet telephony now

Internet-only telephony is actually being used today by consumers. This short-term version of class 3 telephony represents individuals attempting to take advantage of the same price distortions as identified in the class 1 discussion. Using only the Internet, and their personal computers, they establish packet flows and attempt voice communications. [8]

There is considerable skepticism that this short term use of class 3 ITel is commercially significant. The following issues apply:

- The use is attractive only because of current price distortions, which (while entrenched) are not fundamental, and will be challenged by more organized business undertakings.
- The existing lack of support for explicit voice QoS in the Internet makes the quality of the call unpredictable. There is no evidence that the broad consumer market is interested in dealing with these fluctuations in quality.
- Since most consumers do not have a residential Internet service that allows them to be connected at all times, it requires pre-arrangement to receive such a call. This limits the utility of the service.

Without these Internet augments (and the others discussed above), this use of Internet is a bit of a "hobby" application, and there is little evidence that it will grow if all it does is emulate simple POTS calling. In the long run, a pure class 3 offering will not survive as a simple POTS-style telephone replacement, but will be enhanced by new features that advance it away from a POTS equivalent, as described in the speculation above. The simple class 3 ITel will survive only if regulators attempt to suppress the class 1 business, in which case the class 3 variant might persist as a consumer-activist campaign for price reform, keeping pressure on the commercial providers. It must be the case that if this use of Internet starts to become widespread, either it will be regulated to make it relatively less attractive, or other prices will shift. So it seems unlikely that this represents a use with long-term wide-spread benefit, so long as the functionality is simple POTS replacement.

---

[8] There are a number of vendors of PC based ITel software today. For a collection of pointers to product information, see the web site at http://itel.mit.edu.

On the other hand, because this class 3 version of ITel more resembles what a long-term Internet telephone service might be, it can be a platform for early market entrants to position themselves while they gain experience. In the first stage of this form of IT, the major business opportunity is providing software to the end-user. The major supporting service that is required for telephony is a directory service that allows users to locate each other. Providers may position themselves to be major players in a later, more mature version of class 3 services.

A comparison of POTS-compatible telephony -- class 1, 2 and 3

The short term Class 3 option described above, in which computers are being used today to call each other across the Internet , is in its current form essentially a POTS-like service (lower voice quality and usability, but POTS-like in the general nature of the service). This, and a hybrid between the two discussed below, are summarized in table 2.

	Application	Beyond POTS?	End-node fcn?	Ease of use	Internet Augments				
					QoS	Pricing	Reliability	Always on	Ubiquitou
1	LD POTS	no	none	same	no	no	local	no	
2	hybrid	no	none	same	yes	yes	yes	no	
3	I'net POTS	not now	major	better?	yes	yes	yes	yes	

Table 2: Comparison of Class 1 Long Distance POTS, short term Class 3 computer-to-computer ITel, and a Class 2 hybrid.

The short term Class 1 and Class 3 variants are similar in that they are essentially POTs replacements, but they differ in a number of other important characteristics. The class 3 option suffers because it requires the full range of augments to become successful. As discussed above, without QoS support in the Internet, a sufficient base of attached users, and an "always on" attachment model, Internet-based telephony today is somewhat of a "hobby" application.

The class 2 hybrid of POTS-style ITel

The descriptions above of the class 1 and class 3 variants of this service represents end-points on a spectrum. There are a number of places between the two that could be realized. This intermediate, which we categorize as class 2, uses the same gateways between Internet and PSTN as the class 1 option, but extends the service so that users directly connected to the Internet can interconnect to PSTN users. Computers and telephones can interoperate.

One way to assess the importance of this sort of proposal is to ask whether it adds appeal to either of the class 1 or 3 end points. Looking at the class 1 variant, with gateways connecting into the PSTN, there is little intrinsic benefit to extending this so that calls can be completed over the Internet. First, almost anyone with a computer also has a phone, and can be reached without adding this option. In other words, the "network externality" represented by the telephone system swamps the externality represented by the Internet today. Second, there are substantial infrastructure implications if one attempts to reach all the way to the end-user over the Internet, because the full range of augments to the existing Internet infrastructure to support QoS, billing, and "always on" access will eventually be required. For the pure class 1 option, as noted above, to provide long distance or international POTS over Internet, one can build a dedicated infrastructure, which removes the necessity of these augments.

Looking instead at the class 3 variant as a starting point, adding the ability to cross connect to existing PSTN end-points burdens the application with the limitation that it must interwork with POTS-style restrictions, and can never evolve to new forms of service. Once that restriction is accepted, since the telephone end-points so outnumber the Internet end-points, the demand for the service will be generated by the telephone end-points, which is the class 1 situation. We therefore conclude that in the abstract, the class 2 option of simple POTS replacement does not add much vigor to the simpler class 1 option, and imposes a very burdensome restriction on the class 3 option.

However, there is a specific context in which the class 2 option has benefits. Calling in the direction *from* a computer *to* a telephone provides two specific short-term benefits. Since the charges for a long-distance call are normally charged to the sender, calling *from* a computer over the Internet avoids the charges associated with the telephone system, and moves the cost into the Internet context, which is currently flat-rate. At the same time, calling *to* a telephone bypasses the "always on" Internet requirement, and makes it possible to complete a call without prearrangement. Thus, calls from Internet to a telephone have benefit in the short run. Note that this has nothing to do with the power of the computer, but only with the current costs and features in the two regimes. This class 2 hybrid is an excellent example of a short term opportunity with no obvious long-term utility.

The class 2 option of POTS-style interconnection between Internet and PSTN raises interesting business questions. If the Internet were to be connected to the telephone system in a widespread way, it is not clear who would install, operate, and benefit from the gateways. They could be installed by Internet providers, by the telephone companies, or by third parties. The revenue situation is very different, depending on whether the presumed model of calling is from Internet to PSTN or the other direction, and whether the goal is to keep the call in the Internet or in the PSTN for the maximum time. The assumption in most cases is that the Internet will have a lower incremental price for a call, so the goal is to keep the call in the Internet. This implies that the existing telephone service providers will view this service with hostility, and they will need some further motivation (e.g. defensive offense) to deploy Internet phone gateways. Internet service providers might deploy these boxes if they can justify the cost as a part of their "total service offering". This would represent a way to tie the lower level Internet service to the higher-level ITel service, an example of vertical integration in the Internet industry. Third party providers will deploy the gateways only if they can derive revenues, which implies that they must bill someone for the use; this billing will add complexity to the basic service.

From The Present To The Future

The examples and discussion above provide a sufficient context for us to summarize our high level assessment of the ITel arena. Our thesis in this paper is that the goal of the class 1 applications is primarily a cost savings objective, which may not have strong long-term durability as a business opportunity. The class 2 applications will evolve to provide speech access to a wider variety of network resources, and thus diverge from simple POTS service. The class 3 options represent the important long-term outcome of the ITel evolution, but the final speculative form of class 3 Internet telephony described above is not practical today, because the necessary augments to the Internet are not in place, and the utility and usability of the service features to the consumer has not been demonstrated.

It is our belief that Internet telephony will evolve as a series of incremental steps. Early variants of ITel will be identified that can be deployed without requiring all the augments described above to be fully in place. These offerings will serve as experiments to prove the market, evaluate demand, explore the desirability of features, and motivate the fuller deployment of enhanced Internet service. The long-term options will evolve to more advanced forms, as the demand becomes clearer, and the necessary augments come into place.

To further elaborate these points, this paper will describe a number of these short-term variants of ITel, to illustrate the requirements and potential of each in taking us from where we are today to the speculative future described above.

Other class 1 examples:

Application		Beyond POTS?	End-node fcn?	Ease of use	Internet Augments				
					QoS	Pricing	Reliability	Always on	Ubiquitou
1	loop bypss	no	little		same	local	no	local	loc
1	vpn voice	no	none		same?	yes	no	yes	no

Table 3: Two class 1 ITel applications -- local loop bypass using cable, and voice over private infrastructure

Bypass of local loop by cable providers

The goal here is to use Internet infrastructure to deliver telephone service over alternative technology (specifically cable infrastructure), bypassing the local loop and the local telephone provider.

Cable providers currently provide telephone service over cable using customer premise network interface devices dedicated to telephone service. They separately provide Internet service using so-called "cable modems" that provide high speed packet transport over cable. If they could provide telephone service over Internet, they could provide both Internet access and telephone service using one device at the customer premise, which would appear to provide substantial cost advantages.

The cable industry is currently speculating on exactly what forms of telephony to offer over their cable plant in the short run. Offering full "first line" telephone service raises many issues for that industry, ranging from a direct need to address policy and regulatory issues such as funding of universal service to technical issues such as the need to engineer their infrastructure to the level of reliability of the telephone system (for example, to remain operational when the power is off.) A more advantageous short-term alternative might be to offer "second line" telephone service, which need not be as reliable nor as full-functioned. This would allow them to "cream-skim" the telephone business, and steal high-profit offerings from the existing local-loop providers. For example, they could offer access to alternative long-distance service, and take away the resulting access charges that the LEC would receive. They can supply second lines for fax (which need not work when the power is off), and so on.

Differentiating characteristics

As table 3 illustrates, the only augments required for this application are to provide QoS over the cable infrastructure to mix the voice and data traffic. Cable modems currently do not support explicit QoS, but the current approaches to bandwidth allocation for cable modems could be extended to provide this support.

A very significant aspect of this variant is that while it is fully interoperable with the PSTN, and uses the existing telephones of the consumer, it begins a push of function towards the edge of the network. A device would be needed at the customer premise to control the connection of the consumer's telephones and house wiring to the two external phone services -- the cable and the copper loop. This box then begins to take on functions of call control, selection of provider based on cost, and so on. It thus represents a small step toward the migration of telephone function out of the network and into the end-point.

Shared use of private packet networks for voice

Corporations or other users that have procured private Internets, whether built from trunks or switched infrastructure, may attempt to carry some of their voice traffic over this infrastructure in order to make use of this investment. Since network capacity comes in large chunks (T1, DS3, OC3, and so on), there may be economies of scope that derive from combining voice and data over one infrastructure. [9] This application resembles the class 1 long-distance variant of POTS-style telephony, except that the wide area infrastructure being used is one that is operated by a private organization, rather than as a public offering.

Other class 2 examples:

	Application	Beyond POTS?	End-node fcn?	Ease of use	Internet Augments				
					QoS	Pricing	Reliability	Always on	Ubiquitou
2	comp tel	slight	some	better?	no	no	PC:yes	PC:yes	
2	Web voice	slight	yes	better?	yes	yes	yes	no	
2	info access	major	no	better?	no	no	no	no	

Table 4: Three class 2 applications -- computer telephony, adding voice to the Web, and voice access to Web information.

Computer Telephony -- Use of computer to control the telephone

This application is not strictly Internet telephony, but rather computer-mediated telephony. The concept is to connect the computer to the telephone system so that it becomes a more sophisticated user interface for advanced telephony functions. The computer could receive, process and store voice mail, maintain a log of all incoming calls, store catalogs of called numbers, and so on. The telephone could still be used for the actual communication.

In the home, this opportunity raises several interesting and important issues. The typical home today has several phones, perhaps one in almost every room. We will not see this density of computers in the relevant future. So any significant use of the computer as part of residential telephony must be a hybrid that permits the telephone to be used when its convenience outweighs the primitive user interface.

---

[9]This opportunity applies to Internet, and also (and currently more popularly) to lower level switched technologies such as Frame Relay, which provide a slightly better control of QoS.

In the corporate world, many employees have both a telephone and a computer. These are currently managed separately, and (for example) when an employee moves, both must be changed separately. Assuming that the employee will continue to have a computer, the opportunity here is to use the computer and its network infrastructure as a replacement for the telephone and the PBX to which it is attached. This would reduce two systems to one, with presumed cost savings.

Differentiating characteristics

As table 3 summarizes, this application moves functions to the edge of the network (a computer or similar consumer device) while continuing to interwork with the existing PSTN and the resulting POTS-style service. It implies high reliability of the computers and networks, and computers that are always on and available to process calls.

Adding voice to the Web

One of the current emerging examples of voice as a component of multi-media computer application is adding voice communication to Web pages. The concept is that a user browsing a Web page can click a button and talk to a representative of the company providing the Web page, thus merging the Web with "800 numbers".

This is described as a class 2 application because the client side is a computer, while the current implementation at the server side is to connect the incoming call into the existing call dispatching equipment that deals with incoming PSTN calls. This is a rather powerful hybrid, because it could be possible for the representative receiving the call to have access to the computer information that the customer is seeing. So this option mixes POTS-style telephony and computer-based multi-media functions.

This application, if mature, will shift 800 traffic onto the Internet, which could adversely impact the existing telephone providers. The voice calls are carried as far as possible across the Internet, and only connected into the existing telephone system at the premise of the Web provider, so the call provides no revenue to the telephone providers.

Differentiating characteristics

Since this application carries the voice across the Internet as far as possible, connecting into the existing telephone system only at the premise of the server, several of the Internet augments discussed above will be required to make this service real, in particular wide-scale introduction of QoS and the related pricing mechanisms. However, it does not require "always on" operation (the consumer originates the calls) nor does it depend strongly on ubiquitous Internet deployment. Any consumer with the service can fully benefit from it.

Voice access to information on the Web

In the long run, the future of the class 2 hybrid (PSTN/Internet interconnection) is not to provide simple voice communication between humans, but to provide voice access, within the POTS paradigm, to a range of new Internet-based services. The use of touch-tone selection to navigate services is a primitive example of this, but the mature form will involve computerized voice understanding, and conversation between a human on a telephone and a computer, which then reaches out into the Internet to obtain services and information for the user.

Services such as this exploit telephone calls across the PSTN, but do not overall much resemble a classical phone call, since the device at the other end is not a human but a computer, and the goal of the call is to obtain access to information located within the Internet.

Differentiating characteristics

This service requires no augments to the Internet, as the voice only passes over the PSTN. These sorts of applications can be deployed as soon as the speech recognition issues have been resolved.[10]

[10]Early examples of systems that employ computer-based voice understanding are available in the research lab today. For example, the Jupiter system developed by the Spoken Language Systems group at the MIT Lab for Computer Science allows a user to call and have a phone-based conversation with a computer concerning the weather, based on current weather information that the computer retrieves from the Web.

Class 3 examples

Application		Beyond POTS?	End-node fcn?	Ease of use	Internet Augments				
					QoS	Pricing	Reliability	Always on	Ubiquitou
3	teleconf.	yes	major	better?	??	no	local	no	
3	chat, etc.	yes	major	better?	yes	yes	yes	no	
3	telework	yes	major	better?	yes	yes	yes	yes	

Table 5: Three Class 3 applications -- teleconferencing, consumer multi-person applications, and telework

As discussed above, the long-term form of class 3 Internet based telephony involves a major migration away from the assumptions of the traditional telephone system, and requires a number of augments to the current Internet. However, there are a number of class 3 applications that are more limited in the objectives and thus in the augments that they require. These represent first steps that the industry will take, as it explores the space of real Internet-based voice communication.

Teleconferencing

This application exploits the power of the computer to move beyond simple POTS service in support of human communications. The objective is to augment simple voice communications with other modes such as video, multiway (many-to-many) communication patterns, and shared workspace and other groupware tools.

Differentiating characteristics

There are several patterns of teleconferencing. One is within the corporate campus. Achieving the necessary QoS might be achieved by over-provisioning in the short term, which will permit this application to be initiated without explicit QoS support in the Internet routers. The second pattern is among the sites of a campus. This will require support from the Internet technology similar to that discussed above under the heading "Shared use of private packet networks for voice"

The third mode of teleconferencing is telework, which is sufficiently different that it is discussed separately below.

An important aspect of Internet teleconferencing is the many-to-many mode, which is more natural in the Internet, but requires special arrangement (a conference call) in the telephone system. Many to many communications greatly benefits from Internet multicast, which is only now being deployed, and which still suffers from concerns about scale and robustness.

Chat room, games, and other on-line real-time applications

Chat rooms and open group interaction share with teleconferencing the basic objective of linking together a number of people in a common real time experience. However, the details are very different.

First, this application is a leisure activity, and thus targets the consumer at home, not the worker at a place of business. Second, the groupware components might be very different. Third, the open nature of the group will call for a more complex directory and group location service.

This could be considered in the context of a hybrid Internet/PSTN mode (class 2), but is more likely to succeed as a pure class 3 application, where the groupware modalities can be exploited to enhance the experience. The possibility of some user having multi-modal participation while others have only voice seems unappealing.

Differentiating characteristics

This application requires QoS, pricing, reliability and reasonable overall bandwidth to the consumer at the residence.

Telework

Telework is a variant of teleconferencing carried out from the home or other remote location[11]. The objective here is to make teleconferencing available at the residence. While corporations can deploy teleconferencing within and among their business sites, it will be the public ISPs that provide (at least the infrastructure for) telework. The bandwidth and service requirements for telework may somewhat resemble the requirements for chat rooms and other recreational applications. This application will thus provide an opportunity for the consumer-oriented ISPs to play in the business sector, and diversify their revenue base.

Differentiating characteristics

Most corporations seem to desire to roll out a telework option only when the infrastructure is widely in place. Thus, there is a requirement for wide-spread deployment of Internet access.

---

[11]More generally, telework could imply only more primitive network access for shared files, but we choose to use the term here to capture the human to human communications aspects of remote work.

Conclusions -- The Long Term Implications Of ITel

Application		Beyond POTS?	End-node fcn?	Ease of use	Internet Augments				
					QoS	Pricing	Reliability	Always on	Ubiquitou
1	LD POTS	no	none	same	no	no	local	no
1	loop bypss	no	little	same	local	no	local	loca
1	vpn voice	no	none	same?	yes	no	yes	no
2	comp tel	slight	some	better?	no	no	PC:yes	PC:yes
2	Web voice	slight	yes	better?	yes	yes	yes	no
2	info access	major	no	better?	no	no	no	no
3	teleconf.	yes	major	better?	??	no	local	no
3	chat, etc.	yes	major	better?	yes	yes	yes	no
3	telework	yes	major	better?	yes	yes	yes	yes
3	gen'l ITel	yes	major	better?	yes	yes	yes	yes

Table 6: Summary of ITel applications

Implications for timing

Our hypothesis is that ITel will evolve from early offerings that can be assembled today into mature forms that have diverged substantially from simple POTS-style telephony. Looking at table 6, some options require fewer Internet augments to deploy, such as teleconferencing, voice access to information, and computer-mediated telephony. These can happen sooner, and will thus serve to explore the demand, and position players in the market. Teleconferencing, for example, will explore the utility of other modes of communication, such as video and shared workspace. The end-node device supporting the user, whether PC or specialized server, will start to implement a sophisticated user interface to ITel services, and will start to act as the user's agent in implementing key ITel functions. In parallel with these first steps, Internet Service providers will start to implement augments such as QoS and pricing. These will permit the next applications, such as voice over the Web and leisure activities such as games and chat rooms. With sufficient success in the marketplace, there will be enough penetration of Internet service to make the final form of class 3 computer-mediated communication practical.

Implications for industry structure

The speculative final form of computer-mediated human communications implies a substantial change in the structure of the telephone service, and the industry that provides it. The design of the Internet, with open interfaces between different service layers, tends to create an industry structure in which the lower level service providers (the ISPs) do not have a substantial competitive advantage in supplying higher-level services such as Web hosting.[12] This pattern differs from that of the telephone industry, where the provision of physical facilities has been linked to the higher-level telephone service. Despite that history, if telephony (or more generally, our class 3 computer-mediated human communication) were to move to the Internet, there is no reason why it would be immune to this separation of lower and higher level services among different players. In that case, local exchange carriers would be shifted into a role where they provide copper loops, and perhaps provide Internet service over those loops, but have no consequential advantage in providing ITel. Only to the extent that the ISP provides a selection of QoS, and a set of supporting services that exist in the network, does he participate in the implementation of the higher level ITel service.

In fact, ITel ceases to be a single unified service at all. As the end-node computer, owned and provided by the subscriber, becomes part of the telephone context, it will accentuate the fracturing of the service. The implementation of key functions at the end node removes them from the control of the telephone industry. ITel becomes a service built out of software purchased by the consumer, operating over a general (e.g. Internet) communications infrastructure that is independent of application.

Another aspect of the fragmentation of the industry is the manner in which secondary services are supported. The services that have been mentioned here include directory services such as white and yellow pages. Yellow pages, which is essentially advertising, can be expected to become a very competitive business, whose final form is hard to predict. But already Web-based services are beginning to appear, and there is no longer any telephone company monopoly on directory services.

Policy considerations

Faced with the prospect of the "vertical disintegration" of the telephone industry, the incumbent players may reasonably try to create circumstances that forestall or mitigate this potential. One "leading indicator" for the long-term form of ITel is the standards and interfaces that are proposed by major industry players. Several examples are obvious. For example, directory services are just emerging, and there is scope for these standards to either encourage or discourage the interworking of different services as a uniform overall directory system.

---

[12]For a discussion of the overall design of the Internet, and the role of open interfaces in fragmenting the market segments, see chapter 2 of [CSTB 94]. For a discussion of some of the specific business consequences, see [CSTB 96].

The form of ITel described here, with major functions implemented in the end-nodes, can only be realized if the necessary control information (what the telephone industry calls "signaling") is carried end-to-end across the Internet. This is the natural mode for Internet control protocols, but one could imagine an attempt to close or restrict these protocols in some way. If, for example, the called end-node could not determine the number and identity of the calling party, this would help maintain ITel as a centralized application. If a class 2 hybrid form of ITel does succeed in the market, the interconnection of the telephone and Internet telephone service will require some interface to telephone signaling protocols. The two extremes are use of touch-tone signals, or the internal signaling protocols of the telephone system, SS7. These would represent very different modes of access into the existing functions of the PSTN, and might cause "Internet Telephony signaling" to grow in different ways.

Regulation of ITel has been proposed, and may again be proposed in the future, based on the observation that in some forms it is similar to the service of the PSTN.[13] Our hypothesis is that most of the class 1 alternatives in fact exist to exploit pricing distortions, and thus arguments about regulation are germane to class 1 ITel. However, if our hypothesis is true that the final form of ITel will result from the evolution of class 2 and class 3 applications, it seems more difficult to apply a "similar service" criterion. The class 3 application will diverge greatly from the function of existing telephony. The class 2 applications will continue to exploit a "phone call" to access the Internet, but if the person on the telephone is talking to a computer that is reading a web page as the overall service, it seems difficult to describe the collective event as resembling POTS. This suggests that any attempts to regulate ITel, based on looking at the class 1 variants, will quickly unravel in definitional confusion.

One of the consequences of this predicted evolution may be a gradual loss of reliability, dependability and utility of the "telephone service". This trend is already visible as a consequence of divestiture, where some people are already so confused that they are unable to make long distance calls from pay phones, rogue pay phones lure the unwary, bills are unpredictable, there is no accountability for disrupted service, and so on. There could be a possible consumer backlash that demands a more predictable, reliable and usable service. One outcome could be regulation as to the characteristics of a future "telephone service", regulation more akin to safety and consumer protection than pricing and cost recovery.

As the industry becomes more and more fragmented, it will become less and less easy for the various industry player to meet, and discuss and agree on common objectives and approaches.

---

[13]See [Werbach] for a recent discussion of options for regulation of the Internet.

This will contribute to a further deterioration in the overall utility of the service, and a resulting pressure for forces other than those of autonomous private industry to come to bear on the problem. These issues will have to be addressed and resolved in a multi-national context.

Acknowledgment

The author gratefully acknowledges the comments and suggestions of Sharon Gillett, Andrew Sears, Russ Newman, Bill Lehr, John Wroclawski and Lee McKnight. This work was funded by the M.I.T. Internet Telephony Consortium and its industrial partners.

Citations

[Braden] R. Braden, editor, *Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification*. Available on line as ftp://ietf.org/internet-drafts/draft-ietf-rsvp-spec-16.txt

[CSTB 96] The Unpredictable Certainty: Information Infrastructure through 2000. Computer Science and Telecommunications Board, National Research Council, National Academy Press. 1996

[CSTB 94] Realizing the Future: The Internet and Beyond. Computer Science and Telecommunications Board, National Research Council, National Academy Press. 1994

[McKnight&Baily] L. McKnight and J. Bailey, editors. Internet Economics. MIT Press 1997

[Werbach] K. Werbach. *Digital Tornado: The Internet and Telecommunications Policy*. OPP Working Paper 29, FCC.

[Wroclawski] Wroclawski, J. *Specification of the Controlled-Load Network Element Service.* Available on line as ftp://ietf.org/internet-drafts/draft-ietf-intserv-ctrl-load-svc-05.txt.