Active Networks In Telephony

N. F. Maxemchuk
AT&T Labs - Research
Florham Park, N.J.

Abstract - DARPA and the NSF are encouraging research on active router networks. By applying these concepts to telephony we:

- **create a commodity, other than bandwidth, that a telephone network owner can sell,**
- **create new businesses for both equipment manufacturers and for network providers,**
- **provide local access with a distinct advantage over the current local operating companies,**
- **integrate the telephone network with the Internet so that traffic can move between the two,**
- **make it easier to assimilate new technology into the telephone network, and**
- **increase the rate of change in networks by reducing the need for standard protocols.**

To show the effect of active networks, the telephone network is divided into three layers corresponding to the user, the network manager, and the network owner. We give specific examples of the services and devices that are enabled at each layer.

I. Introduction

The current Internet protocols[1] are the result of the layered architectures that were used in the original ARPA Net[2], IBM's SNA, the ISO/OSI seven layered architecture[3], and a number of other data networks. Layered architectures isolate the different processing and transmission technologies in a network, so that different components can change without affecting the whole. This strategy has encouraged advances in technology by controlling the effects of changing different components.

Most layered architectures were designed in the 1970's and are based upon the technology of that period. The most significant change that has occurred since that time is use of processing and memory in almost all aspects of our lives, from automobiles to dishwashers. Increasing amounts of processing are now being used in switching and transmission facilities. For instance, ATM is capable of more complicated sharing and switching functions than the multiplexing strategies of the 1970's.

The isolation provided by layered architectures makes it difficult to make the full range of capabilities in new switching and transmission facilities available to the user. For instance, while it is possible to transmit IP over ATM, it is difficult to make the full range of quality of service in ATM available to user applications. The isolation of layered architectures is beginning to work against us. We should expect that the trend toward intelligent components will continue and that ATM is just the first member of a future family of intelligent transmission systems.

The active networking strategy makes the processing components in the network available to the user[4,5]. The current research on active networks is aimed at data networks, however, the concepts can also be applied to telephone networks. Our objective is to determine how to apply active networks to telephony and to examine the advantages that can be obtained.

According to telephone folk lore, the telephone network is in danger of becoming disabled whenever the processor in a switch is made accessible to the telephone engineers, never mind the average telephone subscriber. A very few years ago computer owners would not have considered making their processor available to unknown users over the Internet, however, that is precisely what most WEB users are doing. The restricted capabilities and controlled environment of Java are proving to be safe. The initial research on active networking is determining safe programming languages and environments.[6].

In a telephone network, active networking is divided into three layers, the user layer, the network manager, and the network owner. At each layer a user program is given access to different network functions. A program introduced by a user can only affect that user, the network manager can affect the set of users that are part of his network, and the network owner has access to all of the capabilities of the network.

In section 2 we will also show that a user layer active network provides a means of interconnecting the Internet and telephone network and also for providing continuous presence to an ISP without tying up home telephone lines. The first capability allows the Internet to bypass parts of the telephone network for packet phones. It also allows the telephone network to bypass parts of the Internet to more quickly deliver WEB pages. The second capability provides more economical local access. Local access companies are being forced to increase the number of switches in their local offices because Internet users and workers who telecommute are increasing the fraction of the time that the typical telephone line is in use. Active networks make it possible for the switch to release the line while the user is idle and automatically reestablish the connection when it is needed. Reducing the holding

times decreases the number of switches that are needed by a local company.

The second layer of active network, discussed in section 3, encompasses the functions that can be performed by a network provider. This layer has fewer restrictions than the user layer. The operations that are allowed increase from the operations that a user performs on his own traffic to the operations that are performed on a group of users. The capabilities that the Java-like environment provides are still restricted because the network manager may not be the network owner. A network manager may purchase lines from a network owner, tailor a virtual private network, and resell a value added service. This layer of active networking can also be used for multicasts with quality of service guarantees.

The third layer of active networking encompasses the functions that are performed by the network owner. The objective is to create an active environment where each node in the network, whether it is a switch or a router, can interpret the instructions and perform a useful operation that gets the information closer to its destination. In section 4 I will show how functions at this layer may make it easier to introduce new technologies into the future telephone network.

For each of the three layers there are two levels of active networking. At the first level, the active components perform processing independent of how the code is delivered to the switch. The code may permanently reside in the switch as part of a user's profile, may be downloaded from the equipment provider's Internet site, or may be rented for the duration of a connection. At the second level the code and data travel through the network as a unit, called a capsule. At this level routes may be modified on the fly and we communicate with successive switches on a path by explicit instructions in the capsule. By reducing the implicit understandings about the operation of switches it becomes easier to introduce new technologies that operate differently.

Capsules are easier to implement in the telephone network than in the Internet because the data units in the telephone network are larger and because the telephone network is a stateful network. The capsules on the internet consist of a packet of data and the code to process that data. The amount of code should be small with respect to the amount of data since the the network should still transmit data efficiently. Therefore, it may be necessary to consider shorthand notations rather than carrying the actual code in a capsule. In the telephone network the capsule is the samples in a voice connection and the code to control the switch. Since the amount of data is larger, the amount of code is not as restricted. The samples in the voice data unit are not contiguous, but telephone switches retain state to associate the samples in a connection.

II. The user layer

There are two important differences between the Internet and telephone networks that make active networks more useful at the user level in the telephone network. The additional user services that are made possible are sufficient to justify introducing active components into the telephone network.

The first difference is that an end user on the Internet has a general purpose computer that can perform a wide range of functions, while the end user on the telephone network has a relatively simple telephone with a primitive user interface. Adding intelligence in the telephone network increases the capabilities of the user much more than it does in the Internet.

The second difference is that the telephone network already has a need for multiparty, real-time collaboration, conference calls, while the need for multimedia, computer-aided collaborations over the Internet has not developed. There are many ways that intelligence in the network can be used in conference calls. For instance:

— Security codes are used to restrict access to a conference call. In an active network an individual can prove his identity to a local agent and have the trust transferred, instead of handing out magic sequences of numbers for each conference.
— The number of participants in a conference is restricted by the ability to add voice signals together and the amount of ingress noise that can be tolerated. In active networks there are many ways of selecting or giving dominance to the participants so that the number of participants can be larger.
— The number of receivers in an audio broadcast is restricted because the call is usually configured as a star with a limited number of ports. Larger numbers of receivers can be accommodated by using active components as splitters at intermediate switches.

In the remainder of this section I will examine three user level applications of active networks: voice coding, continuous presence on a ISP and Internet telephony. These three applications show how active networks can:

1. allow the technology used in the telephone network to evolve,
2. provide an advantage in the local access market, and,
3. move traffic between the Internet and telephone network, depending upon which provides the most economical service with a specified quality, at a particular time.

A. Voice Coders

The 64 kbps, 3 kHz, voice coder is the poster child for technology stagnation in the telephone network. When the coders were developed, in the late 60's, early 70's, nand gate cost 25 cents a piece and J-K flip flops were a buck a bit. The 64 kbps voice coder was as complicated as was reasonable. In the last quarter of a century memory and logic prices have decreased by many orders of magnitude and speech coding algorithms have improved. It is reasonable to provide the same quality signal with less than a quarter as many bps or to provide a higher quality signal, adequate for speaker phones, in the same bit rate channel.

The problem is that the 64 kbps coders are hard wired into the network and it is difficult to change some without changing a large number simultaneously. The coders are hard wired into the network both literally, in that special purpose circuits perform the coding, and figuratively, in that the entire network revolves around the magic number 64,000.

Active networks provide a mechanism for modernizing network coders without disrupting the entire network. The local loop, from the end user to the first switch, can support much more than a 64 kbps digital signal or the 3 kHz analog bandwidth that is used for current telephone calls. Therefore, the quality of a connection can be improved by changing the operation of the switching network without changing the local loop. Active networks make it possible to load software for a different coder into the switch for each connection. Therefore, each connection can select a different bit rate or bandwidth. We can also select the current 64 kbps coder whenever communicating with an end user on the inactive part of a network. Therefore, the advantages of active networks can be introduced gradually.

End equipment manufacturers can distinguish themselves by providing better speech coding algorithms and may lease the coders by the connection. Software rental is reasonable in an active telephone network because the software is used in the network and is not available to the end user, so there is less concern with illegal copying. Long distance providers can cut the cost of their service by translating between 64 kbps and a lower rate coder, perhaps 16 kbps, and using the lower rate over long distance lines.

B. Continuous Presense

The precipitous rise of the Internet has created a problem for the local operating companies. The switches that service the lines in residential areas are designed to support about 6 ccs (600 call seconds per hour). The number of switches in local telephone offices were selected with the assumption that, even during the busiest times of the day, only $(1/6)^{th}$ of the lines are being used. Internet users are connected to the network longer than telephone users. As a result, more lines are simultaneously in use and local telephone companies are being forced to increase their investment in switches.

The effect of the Internet on the telephone company is just starting. Instead of customers holding for a longer time in the future, it is likely that they will remain connected continuously. Continuous presence has resulted in a wider acceptance and use of email in the office than in the home. Fewer people use email when they have to check mail boxes than when it automatically appears on their computer screen. In a residential environment, continuous presence will also create a new generation of "push-based" services, such as electronic delivery of periodicals, reminder services, and aggressive advertising.

Continuous presence will cause a greater increase in the investment in local office switches. The problem will be compounded as more office workers telecommute. With active networks every user can be provided with continuous presence to an ISP and his office with a smaller increase in the number of switches. A program in the switch can detect idle connections through the switch, tear them down, but remember the destination and reestablish the connection when the user becomes active. There are a number of ways to decide when to make or break a connection. The advantage of active networks over a hardware approach, or even a centralized software approach, is that not every user needs to use the same technique, so that the techniques can be opened to competition.

C. Internet Telephony

A year or two ago Internet telephony threatened to invade the domain of the telephone company by providing a low cost, or no additional cost, alternative to long distance and international calls. The threat has not been great because the voice quality has not been dependable.

Once the quality concerns are solved, Internet telephony has advantages beyond price. The PC that is attached to the Internet is a much more powerful user interface than the touch tone pad on a telephone. Therefore it is possible to offer a richer selection of services, such as more elaborate teleconferencing with more elaborate controls for adding and deleting users or directing remarks to subsets of the participants, scheduling and initiating voice meetings by email, etc.. It is also possible to change voice coders and to use different quality coders for telephones, speaker phones or simple activity monitors. The advantages that become

possible in a user level active telephone network can already be achieved in the Internet.

Active networks can use the telephone network to improve the quality of the Internet, so that the Internet can carry good quality voice. A capsule can carry the program and data that are needed to monitor the delay and loss as the capsule moves through the network. When part of the Internet becomes congested, an active node can place a phone call and use the circuit switched telephone network to bypass the congested segment. A bypass strategy would search for the least cost combination of Internet and telephone links that provide the required quality of service.

III. The network manager layer

The network management layer deals with a group of users, rather than a single user. The instructions that can be performed and the facilities that can be accessed are less restricted than those at the user layer. The most obvious example is the use of active networks for access and flow control. An individual user can control his access rate and decide which of his packets to discard. He cannot decide if his packets are more important than those of another user. A network manager, on the other hand, can provide more delay, but lower losses, for data than voice or video by making decisions across a community of users.

In the remainder of this section we will consider three possible applications at this layer. The first is virtual private networks for telecommuting, the second is multicast video and audio, and the third is bypass networks. The first application is an improvement on a service that telephone companies currently provide. The second application is a service that the telephone network should take over from the Internet, since the long duration and continuous transmission is better serviced by circuits than packets. The third application shows how circuit and packet networks can be combined to improve the cost and performance of both.

A. Virtual Private Networks

The long distance and local telephone companies and communication management companies configure special purpose private networks for corporations. Trust and security within the private network is generally greater than outside, since the others on the network are presumably part of the same company. The networks may use fire walls to provide a greater degree of security for data terminals on the VPN than terminals that are connected directly to the Internet. Telephones on the network may have access to groups of lines that interconnect locations, and calls between individuals inside the company usually require dialing fewer digits.

A challenge in a virtual private network is to add and delete lines from the public dial-up network to provide telecommuters and small satellite offices with the same services and security as individuals in the central offices. Active networks can make the appendages appear as part of the central network, even if the composition of the network changes frequently. An active network can authenticate individuals and then provide them with the same protections and privileges as the rest of the network. Integrating the public and private networks in this manner also makes it possible to return lines from the VPN at the end of the work day.

Determining the range and type of services in this type of a virtual private network is particularly important because a long distance carrier's network, or a subset of that network, is a special case of a VPN. A long distance carrier can provide directory services, or the ability to use compressed links to customers who are part of its VPN. Similarly, a long distance carrier can set up different grades of service as a VPN that consists of a subset of its own customers. For instance, some customers may want access to shared T1 or T3 links between distant locations, in the same way that an individual on a corporate network has access to the interfacility links.

B. Reclaiming The Mbone From The Internet

On the Internet, the multicast backbone (mbone) is used to distribute real time sessions from most technical conferences and IETF meetings. The system uses PC's and workstations to compress the audio and video signals, rather than the special purpose hardware that was used in earlier telephone network trials and services. As a result there is a very large potential audience.

The Mbone can also cover news events that are of interest to a large number of people, although not a large enough concentration to warrant broadcast coverage. For instance, when there is an earthquake, emigrants from that area want much more detailed coverage than the general population. There may be many emigrants in different parts of the world without a major concentration in a specific city. Other applications might include university courses and certain entertainment events.

The problem with the mbone is that the quality of the signal cannot be guaranteed. We recently measured the packet loss rate in Holmdel for a seminar series transmitted from Berkeley. Parts of the broadcasts were clear, but parts had loss rates exceeding 50% and were useless.

Real time video and audio, that is compressed for transmission on the Internet, transmits high, but time varying packet rates for a long period of time. This requirement almost matches the characteristics of a circuit

switch. With active networking the match can be made perfect. The circuit can operate at the upper bound of the coder's rate and return unused bits for packet transmission.

The multicast applications on the Internet require the quality of service of the telephone network in order to prosper. Once the quality is improved there should be a large pay per view business, where entrepreneurs lease a multicast tree from the telephone network to distribute special events.[7]

C. Bypass Networks

Packet networks provide a means of sharing network capacity while frequently and quickly changing connections. Circuit switched networks provide a means of dedicating capacity for guaranteed quality, long lived connections.

An ISP can use the circuit switched network to reduce the delay that its users are experiencing across congested regions of the Internet. When the delay or loss across a region of the Internet exceeds the end users patience, the ISP can add a dedicated circuit so that its customers do not travel through the congested region. Competition on the Internet is as fierce as on the telephone network. Therefore, it is reasonable to assume that if one ISP noticeably reduces the "world wide wait" then they all must.

The long distance telephone company is not necessarily the company that owns the circuit switched facilities. A long distance telephone provider can use uncongested regions of the packet network to bypass parts of the circuit switched network. The long distance provider can lower its own cost by using the Internet to replace dedicated circuits where and when the Internet can provide reasonable quality at a lower cost.

4. THE NETWORK OWNER LAYER

The final layer is the network owner. The network owner has access to all of the network's resources, and there are no restrictions on the programs that the network owner may place in the switch. The principal advantage of active networks is that it makes it easier to incorporate new technologies into the network so that the network can evolve continuously rather than suffering disruptions when its technology becomes obsolete.

The last major disruption in the telephone network was the change from analog to digital technology, which was delayed at least 10 years because of the embedded investment in reed relays and other analog devices. At present, ATM technology has been delayed at least 10 years because of the investment in bit multiplexers. Active networks make it possible to introduce new technology into small parts of the network continuously.

There are two levels at which active networks enable technology evolution. At the first level there is a common programming language that reduces the need for implicit agreements between switches The first level of active networks allows switches from different vendors, with different capabilities to be plugged together. The program is a more powerful connection, that can accommodate a wider variety of operations, than a common transmission protocol.

A common programming language also reduces the need for most standard protocols. For instance, in today's environment it would be natural to standardize the format of a capsule, however, in an active network we don't even have to specify how a capsule ends. The program at the beginning of the capsule can instruct the switch. If the data is an IP packet, the length is know when the capsule is formed and can be specified in the program. If the data is an active interval in a telephone conversation, the length is not known at the beginning of transmission and a data link escape sequence, such as those used in IBM's SDLC and BSC protocols, can indicate the end of capsule. Different capsules on the same link can use different techniques, and there is no need for a standard protocol. When we eliminate standards, we also eliminate the time needed to approve new standards and increase the rate of change of the network.

The first level requires that the switches use the same transmission format so that the links from one can be plugged into the other. This level is not sufficient to connect a TDM telephone switch directly to an Internet router or even an ATM cell switch. Greater flexibility can be obtained by isolating the transmission and service layers in switching.

At the second level of active networking, the capsule is used as the common component that connects different transmission technologies. The program at the beginning of the capsule provides instructions for an intelligent component to translate the data in the capsule into the format that is expected by the switch or router. The use of capsules in an active network is similar to the use of the IP packets in the Internet. The IP protocol isolates services from the transmission facilities. The Internet transmits IP packets over everything, while TCP and other services operate over IP, independent of the transmission facilities. The difference between an IP packet and a capsule is that the capsule can better use intelligent components.

A. Switching Language

The objective of the switching language is to replace the rigid communications between network components with a more flexible mechanism. At present, successive

components use standard protocols to communicate and establish and maintain a path between a source and destination. The standards are difficult to agree upon and change, so new components try to operate within the framework of the old standards, constraining the evolution of components. The switching language is much more general and provides explicit instructions to the successive components. The program may implement a protocol that is currently a standard. Since the program that controls communications must satisfy the live lock, dead lock, and other safety properties of protocols it is very likely that a small set of protocols will be used for most communications. The difference between the protocol that is implemented in the switching language and the current operations is that the protocol does not need to be standardized or even used by the other communications that are sharing the switching component.

Tennenhouse[4] described the user level language in an active network as an extension of the technology that led to Postscript and Java. The safety properties of JAVA have made the concept of users programming a network acceptable. There were pre-Java attempts at network programming languages[8], but they lacked the ability to constrain the damage that could be done to the networks operation. At the network owner layer I look at the language as being more closely related to HTML than JAVA. HTML is a set of instructions on the general, rather than specific, appearance of a document. Each device that the document is displayed on may have a different window size or set of fonts. The device uses the instructions to present a reasonable document, although different devices will create documents that look very different.

The flexible interpretation of HTML is more important at the network owner level than the safety properties of JAVA. The switching language should have general, rather than specific, instructions that specify bandwidth, quality, delay, etc. Different devices (a TDM switch, an ATM switch, a packet router, etc.) may interpret the instructions differently in order to provide a reasonable communication path. Even within a class of devices, different manufacturers may interpret the instructions differently to distinguish their products. By installing different sets of switching devices, telephone companies can provide different looking services, at different costs, to meet the same general requirements.

No matter how carefully we select the language that is used in active networks, it will have to change to meet new challenges. The evolutionary path for capsules is the same as the evolutionary path for IP. The next generation capsule will be encapsulated in the current generation capsule in order to tunnel between islands with the new capabilities. The intelligence in a capsule adds capabilities to "capsules within capsules" that exceed those of "IP over IP." For instance, the program in a capsule may perform a translation between the new and old capabilities in order to connect a user on one of the islands to the vast majority of users in the sea between the islands. On the Internet, IP over IP is used to connect users on the islands that have the same capabilities.

B. Transmission Layer

The objective of the transmission layer is to separate transmission facilities from the services that are provided by the telephone companies so that one can change without affecting the other. This is similar to the IP layer in the Internet. The IP layer is designed to operate over a number of transmission facilities, (ie. FDM, TDM, WDM, 50 kbps, T1, T3, OC3, ...) so that the facilities can be mixed together in the network.

In contrast to the Internet, the telephone network is "vertically optimized". There is no isolation between layers. Any layer can use knowledge about the other layers to reduce the cost of performing a function. For instance, a switch in the TDM network uses the time between slots and the number of connections in the switch to constrain the services that the switch can implement. Vertical optimization makes the best use of existing facilities and leads to the least cost solution for a particular technology, but makes it very difficult to change any part of the network.

The Internet uses current technologies less efficiently than it might in order to be better able to change as new technologies evolve. In effect the network is more expensive today than if it optimized the use of current technologies, but will be less expensive tomorrow than if it did not incorporate a new technology. This strategy may not be reasonable in the competitive environment and low margins of the telephone business. It may be comforting to know that our network will cost less in the future than our competitors, but if we have to charge more today, we may not survive long enough for it to matter.

Active capsules are more flexible than standard protocols and provide the means for both optimizing the use of the current technology and changing in the future. The program in a capsule contains instructions that can be implemented differently in switches that use different technologies. If a technology change in a network component makes it impossible to reinterpret the commands in a capsule, the previous network component can substitute new commands or encapsulate the capsule in a new capsule that can be processed by the new component. In any event, we can optimize our current combination of technologies and change the programs in capsules when the mix changes.

Capsules provide the means to use intelligent network components. The current IP protocol is designed to limit, as much as possible, communications between services and transmission facilities. This strategy works well over "dumb" transmission facilities, such as dedicated circuits, but cannot take advantage of the full range of capabilities of more flexible facilities, such as ATM. The program in a capsule can transfer instructions to properly use the full capabilities of the transmission facilities

For instance, when IP is transmitted over ATM, IP isolates the service layer from the ATM capabilities. The ATM cells are treated like a fixed rate connection and the sharing and QoS capabilities of ATM are not accessible by the service. A capsule can include a request for a level of service. A TDM switch may provide a constant rate circuit while an ATM switch may use a QoS that is adequate but uses less network bandwidth. The assumption is that a new transmission technology will execute the program to provide an acceptable, although not necessarily identical service.

IV. Conclusion

The ARPA net, which evolved into the Internet, used a packet switching technology that was a deviation from the path that the communications companies were following at that time. The trend in communications was to use regimented time or frequency sharing techniques. DARPA funded the research on statistically shared networks.

Since funding the ARPA net, the initiatives from DARPA have supported the direction of the communications industry. For instance, the fiber optic and ATM switching initiatives had participation by many of the major communications suppliers.

Active networks represent another deviation from the path. The industry trend is to move intelligence out of the network as programming the network becomes frightening, if not impossible. Active networks move more intelligence into the network and use evolving "safe" programming paradigms, like JAVA, to make programming the network less dangerous.

The basis for active networks seems to be solid. The number of university researchers who are working in this area is increasing. Therefore, it is likely that active networks will have an effect on the future communications infrastructure.

The university work on active networks is directed at packet networks. In this note, I explain why the effect of active networks on the telephone network may be even greater than the effect on the Internet.

References

[1] D. D. Clark, "The design philosophy of the DARPA internet protocols," ACM SIGCOMM ’88, Stanford Calif, Aug. 1988, pp. 106-114.

[2] D. C. Walden, "Experiences in Building, Operating and Using the ARPA Network," Second USA-Japan Computer Conference, Tokyo, Aug. 1975.

[3] J. Moulton, "High Level Protocol Boundaries in the ISO Model," IEEE 1980, pp. 54-58.

[4] D. L. Tennenhouse, D. J. Wetherall, "Towards an active network architecture," Computer Communications Review, Apr. 1996, vol. 28, no. 2, pgs. 5-18

[5] U. Legedza, D. Wetherall, J. Guttag, "Improving the performance of distributed applications using active networks," Proc. IEEE INFOCOM’98, San Francisco, CA, Mar. 31, Apr. 2, 1998. p.590-9.

[6] C. A. Gunter, S. M. Nettles, J. M. Smith, "The Switch Ware Active Network Architecture," IEEE Network Magazine, May 1998, pp 29-36.

[7] N. F. Maxemchuk, "Video Distribution on Multicast Networks," IEEE JSAC, April 1997, vol.15, no.3, pp. 357-372.

[8] J. Zander, R. Forchheimer, "The Softnet Project: A Retrospective," Eurocon 88, June 13-17 1988, Stockholm Sweden, pp. 343-345.