Multiplanar applications and multimodal networks

S. Keshav
Cornell Network Research Group
Department of Computer Science, Cornell University
Ithaca, NY 14850
skeshav@cs.cornell.edu

Abstract

We believe that a broad class of future applications will span both the Internet and the telephone network because such *multiplanar* applications have several economic and architectural advantages over conventional ones. We also envision the close interlinking of the telephone network and the Internet to form a *multimodal* network. In this paper, we describe these applications and networks, outline their architecture, and present our experiences in constructing a prototype multiplanar application.

Keywords: Future applications, network architecture

1. Introduction

Computer applications, over the last fifty years, have evolved from programs running on a single mainframe, accessible only to users physically present at the mainframe installation, to distributed, multi-component systems whose parts communicate over a network. We can represent this progress with the following geometric metaphor. Programs running on a single machine are accessible only at single zero-dimensional 'points'. For example, a program calculating ballistic trajectories on a mainframe receives its input and produces output at a single location. If we now introduce an underlying network, we can imagine 'linear' one-dimensional applications that require communication between two 'points' or hosts. This encompasses the client-server paradigm, where a single client and a single server cooperate to provide some service. An example of a linear application is a file server, where a stub program on a client communicates with a server to provide file access.

The next evolutionary stage is a 'planar' application, where the application requires simultaneous communication between multiple hosts. Two examples illustrate such applications. First, consider a cluster-based computation, where an application partitions computation among a cluster of machines. Distributed ray-tracing is a good example of this type of planar application. A second sort of planar application is a component-based application, such as a distributed database, where the application's functionality is distributed among a set of communicating components.

Continuing this progression, the next obvious step is a three-dimensional application. What would the third dimension imply? Recall that the 'plane' in a planar application is formed by a communication network. Thus, in a three-dimensional application, one could imagine that each host would simultaneously be present on a several networks, allowing components to communicate using any of these networks. While such applications may become widely available in the future, the focus in this paper is a more restricted set of *multiplanar* applications: applications that span multiple components, and the different networking planes are linked by *gateways*. Thus, in a multiplanar application, every host can communicate with every other host present in any of the constituent networks, but only with the help of an intervening gateway (Figure 1).

We claim that such multiplanar applications are interesting not just in terms of their position in this evolutionary continuum, but also because they have significant economic and technical advantages over existing point, linear, and planar applications. We discuss these advantages further in Section 2.

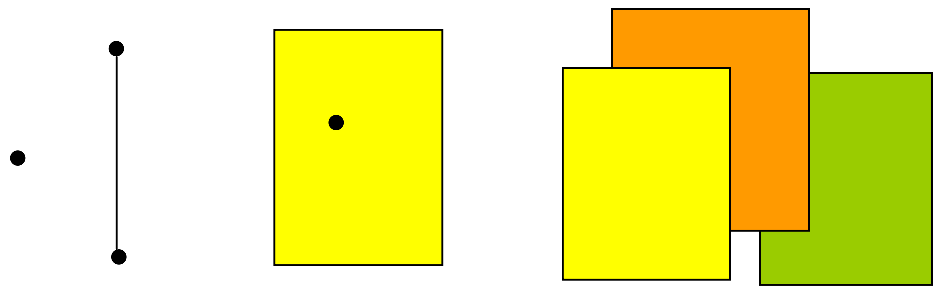

Along with the evolution of applications described above, there has been a parallel evolution in internetworking that we describe next. The world has three major internetworks[1]: the telephone network, the Internet, and television and radio broadcast networks. Each network is independently being extended in *breadth* and *depth*. By an increase in breadth, we mean that each network is covering greater and greater geographical areas, and within each area, is achieving greater endpoint density. Cellular telephones, cable TV, and modem access to Internet are some of the technology drivers behind this trend. The three networks are also increasing in depth, that is, in the number of services available to an endpoint. Pay-per-view service in the television network, fax service in the telephone network, and Web-based services in the Internet are examples of increases in network depth.

While the telephone network and the television network have their advantages, such as guaranteed quality of service, and in the case of broadcast networks, a very low cost per endpoint, they are both much less flexible than the Internet in the creation of new services. Moreover, as the intrinsic bandwidth available in the Internet increases year after year, it is more and more likely to become the dominant network in the near future. This prompts the evolutionary trend of *replacement*, that is, a trend towards using the Internet to replace capabilities formerly provided by one of the other networks. For example, much of the revenue acquired by telephone companies from 800-number service is being replaced by catalog ordering over the Internet. Similarly, with real-time audio and video streaming, some of the functionality of the radio and TV networks is already being provided by the Internet.

Extrapolating from the above trends, one may come to the conclusion that eventually, the entire functionality of the telephone, television, and radio networks will be subsumed by the Internet. While this scenario has its merits, we believe that the existing installation base for these networks is so large, and the revenue that they already generate so substantial, that, at least in the near term, we will see the rapid deployment of *multimode* networks. These are networks built from multiple networking technologies, with interworking between disparate components provided by gateways. Such multimode networks allow users on one network to access resources on another network, thus leveraging the existing subscriber base. Clearly, such multimode networks are the substrate underlying the multiplanar applications described earlier.

We can now tie these two evolutionary trends together. On the one hand, we believe that component-based planar applications will eventually span multiple networking planes, allowing components on heterogeneous networks to interoperate to form multiplanar applications. On the other hand, the lack of flexibility in the telephone network and broadcast networks, coupled with their large subscriber bases, will make multimode networks more attractive. Together, multiplanar applications on multimodal networks are likely to be a significant new class of applications in the near future. Note that a multiplanar application

---

[1] By an internetwork, we mean a network of networks. By definition such networks span multiple underlying (level two) technologies.

cannot exist without a multimodal network. Thus, in the rest of this paper, we will only describe multiplanar applications, implicitly implying that the underlying network is multimodal.

2. Advantages of multiplanar applications

The primary characteristic of a multiplanar application is that it has components operating in both the data and the control plane in more than one internetwork. A concrete example of such an application will help to fix ideas. Consider the following 'Internet voice mail' application. The application allows a telephone user to register an email address with it. If the user receives a telephone call, and does not answer this call for some period of time, then the application answers on the user's behalf, plays a customized greeting, collects the incoming voice message, digitizes it, and sends it to the user as an email attachment. The application has a presence in the data plane of the telephone network, in order to play the greeting, and to collect the message. It is also present in the Internet data plane, because it sends an email message on the Internet. It is present in the telephone network control plane, because it can monitor how long a call is unanswered, and can pick up a call on a user's behalf. And finally, it is in the Internet control plane, because it can initiate a TCP connection with the user's mail daemon in order to transfer email. We now turn our attention to the advantages of multiplanar applications over single-plane applications.

The advantages of a multiplanar application can be summed up as: it provides *new services* to *existing subscribers* using a *well-known user interface*. Let us consider each point in turn. First, a multiplanar application provides services that can be provided in no other way. Continuing with our example, a uniplanar application cannot provide the functionality of an answering machine and email. This ability to create new service offerings is key to the success of multiplanar applications.

The second point is that multiplanar applications are immediately available to existing users on both networks. Any telephone subscriber can make a call into the Internet voice mail application, and the system can send email to any user on the Internet. Thus, once an association between a number and an email address is set up, and calls to the user are directed to the appropriate control element in the application, the application can transparently serve users on both networks.

Finally, because multiplanar applications reuse existing data and control plane elements in underlying networks, they have the same user interface as existing applications, making them easy to use. For example, with Internet voice mail, the person leaving the message has exactly the same interface as an answering machine, and the person retrieving the message uses the familiar email interface. Thus, the new application does not require users to be retrained.

These three features make multiplanar applications economically attractive to service providers. Once installed, existing customers, using existing user interfaces, can access them to obtain a wide variety of new services. Multiplanar applications allow service providers to rapidly create new value-added services. They are of particular interest to telephony service providers, who can leverage their customer base and existing investment in signaling and billing mechanisms to create new Internet-accessible services. Internet voice mail, for example, would certainly be a useful revenue-bearing service for the 60 million telephone service customers in the US who are also on the Internet!

3. Requirements

Multiplanar applications pose some non-trivial requirements on the underlying internetworks. First, they require gateways both in the data and the control plane to allow communication between the components. Such gateways could be transparent, allowing direct communication between components, or may be visible, requiring explicit communication between components and the gateway. An example of a transparent gateway in the control plane is provided by the JTAPI (Java Telephony API) that wraps the capabilities of a telephony PBX in a Java interface. This API allows Java applets to communicate directly with the SS7 stack on the PBX, thus allowing calls to be initiated or monitored from the applet. In contrast, the data gateway between the Internet and the telephone network is typically visible: packets from the Internet must be sent to a specific gateway (which the communicating host must be aware of) in order for these packets to enter the telephone network. Of course, with additional layers of software and appropriate redirection, any visible service can always be converted to a transparent one. Nevertheless, the fact remains that all communication between components on separate networks must pass through gateways. To prevent these from becoming single points of failure, we expect that gateways would be actually implemented as a fail-safe process group, with automatic cutover from a failing primary to a backup.

Gateways, in addition to simply connecting components, may also perform a variety of non-trivial functions. For example, in the data plane, they could provide format conversion in the form of compression or adding redundancy for error correction. In the control plane, they could perform protocol translation, monitor specific network components for failure, or export a management interface to provide control actions in both connected internetworks.

To obtain reasonable communication between data-plane components on different internetworks, multiplanar applications require that the internetworks and gateways provide comparable quality of service. To take a trivial example, a multiplanar application that requires interactive voice communication would probably not run on a Internet WAN link. Thus, when deploying the application, service providers must take care to match the QoS availability with the application's requirements. Note that the Internet provides perfectly adequate service for carrying one-way voice, and, to some extent, one-way video. Moreover, if delay is not a consideration, voice messages can be sent with an arbitrary degree of fidelity on the Internet. Therefore, despite the absence of good service quality on the Internet, it can still be effectively coupled with the telephone or television/radio broadcast network. For example, one can easily imagine a service that allows users to email short video clips to a community access cable channel for local TV broadcast. Indeed, this might be an innovative way for members of a community to quickly cover events of local interest with their neighbors!

4. Examples

In this section we present some illustrative examples of multiplanar applications. While these only paper designs, in Section 5, we will present the design and implementation of a multiplanar application that we implemented in a testbed at Cornell University.

4.1 "The network is the *calculator*"

Sun Microsystems has evangelized the slogan 'The network is the computer'. In this toy application, we use a multimode network to emulate a calculator. The idea is to use the touchtone keypad on a telephone as the input to a calculator. The user dials a telephone number to reach the calculator, and, when prompted, enters two numbers to be multiplied, separated by the '*' key. In response, a window pops up on the screen with the product of the two numbers. The application demonstrates the ability of the multiplanar application to (a) answer a telephone call, (b) learn the caller's telephone number, (c) acquire DTMF (touchtone) digits from the keypad, (d) carry out the multiplication, then (e) invoke the X server on the caller's machine (identified through a lookup on a database with the key as the caller's number) to pop up a window with the answer. While the objective of the application is trivial, the components necessary to carry out this computation are not.

4.2 User-controlled answering schedule

Telephone users often need to have their telephone answered according to a set of choices. For instance, they would like the office number to ring, say five times, and if there is no answer, then to have the number ring in a secretarial pool, and if that fails, to transfer the call to a cell phone or a pager. The choices may depend on the time of the day—for instance, the first choice at night might be the home number, or an answering service. Unfortunately, the interaction between the user and the telephone network to program these choices is poor. Moreover, it is hard to alter the choices, once they have been made. The proposed multiplanar application has one component running in the telephone network, monitoring all calls made to the user's number. On receiving notification of an incoming call, a simple script is run, that decided where the call should be routed. The innovation is to allow the script to be entered either through a Web-based form, or through email. So, for instance, a user might mail the application a set of answering rules from a remote location. This is particularly convenient for users who are away from their home location. The application requires components in the telephone network's control plane, to intercept and divert a call. It also requires data and control components in the Internet plane to receive a script and validate the author of the script.

4.3 Two-way paging

Pagers use broadcast on a well-known frequency to allow a *headend* to communicate short messages to a set of pagers. Web-based paging and email-based paging allow users to use the Web or Email to send a message to the headend; this message is then broadcast and received by the chosen pager. This already is a primitive form of multiplanar application (trivial, because there is almost no control in the broadcast network). 2-way paging, however, is a much more interesting example. With two-way paging, the paging device can originate messages that are transmitted to the nearest receiving base station. The translates the incoming message either to email or to a Web page. The originator of the message can then read email or browse the Web to obtain a reply. The innovation here is to use the Internet to carry messages from the pager back to the page-initiator. This would be impossible to do using only the broadcast network--the Internet provides exactly the one-to-one capability required for the reply.

4.4 Micro-movies

Commercial movie theaters these days screen movies that the theater owners think will draw viewers. They actually have no way to know if a movie will do poorly. What's worse, usually even a poorly attended movie must be shown for at least a week, since prints are circulated only once a week. In this application, we first free the theater owner from the tyranny of the master print, and then make it possible to customize movie showings to the audience. To achieve the first part, we envisage the digitization of movies in some digital format, and the storage of moves either on a central or distributed file repository. Suppose each movie takes 20GB of storage. This can be downloaded over a dedicated T3 link in less than one hour. Given sufficient capacity in the telephone network, and high-speed Digital Access Cross Connects, setting up short-term T3s is technically feasible even with current technology (it would be trivial with ATM). Thus, we claim that a theater could, one hour before show time, decide which movie to show. The choice of movies could be made by letting moviegoers vote on what movie they want. They could do this, for instance, by registering their vote (and credit card number) on the theater's web site. This would allow the theater-owner to automatically gauge the popularity of a movie before screening it! While this example is somewhat far fetched, we believe it shows the types of novel applications that are possible by combining the fast dedicated data-pipes in the telephone network with signaling and flexible distributed control in the Internet.

5. Detailed example – Internet voice mail

We now present the design and implementation of a multiplanar application that we implemented in a testbed in Cornell University. The primary goal of the application is to provide users with convenient access to their voice mail messages by reusing the existing infrastructure for email and audio streaming. Specifically, users are given 4-digit extensions to a shared telephone number. A caller on the telephone number dials the access number followed by the extension hears a customized greeting from the called user. The caller's message is then digitized and stored as a RealAudio file in a Web server. The called user is then sent email containing the URL of this file. By clicking on this URL, the user brings up a RealAudio client that streams the message over the Internet to the user's workstation, playing the message. The system supports two other interfaces. One allows users to record a greeting by calling a special telephone number.

The other is a management interface that lets an administrator assign users with passwords, and associates extensions with email addresses. We now discuss the components of this application in greater detail.

5.1 Data plane

This application has data plane components both in the telephone network and the Internet (Figure 2). In the telephone network, voice samples travel across the telephone network to the experimental PBX. They are then diverted to a special line that is connected to a voice-data modem on a PC.

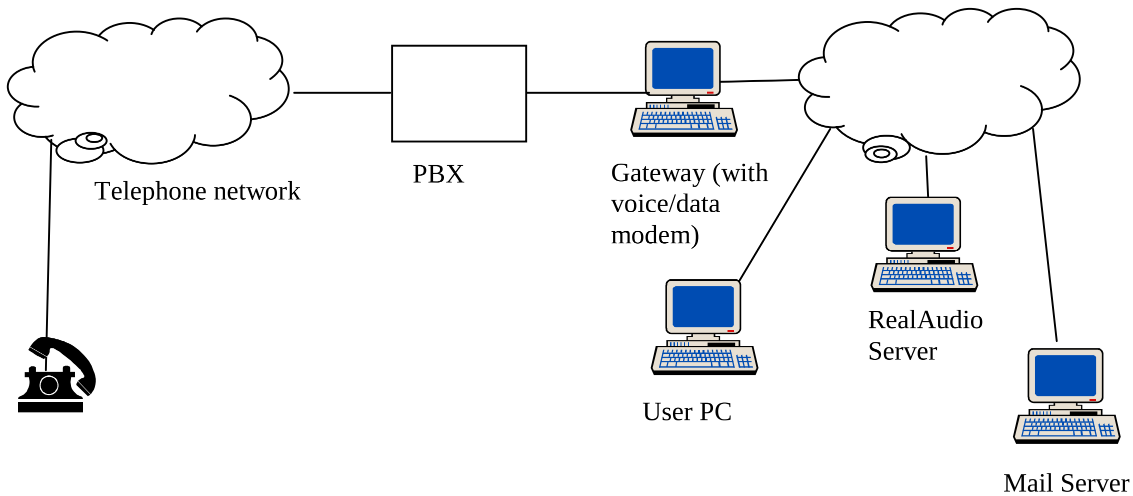

Figure 2: Data plane in the Internet voice mail application

The gateway digitizes voice samples to the RealAudio format, and stores them in a RealAudio server. It also plays a customized greeting on the phone line. When the message completes, the gateway sends mail containing a pointer to the RealAudio file to the mail server. When the user checks mail on the mail server, he or she receives this pointer, and can then retrieve voice mail streamed across the Internet.

When administering the greeting, the user dials a special number, which the PBX diverts to the modem line connected to the gateway. The incoming greeting message is stored as a voice file on the gateway (in WAV format), associated with the extension of the user. This is the file that is played when the user's extension is called. Users who do not register a customized greeting are assigned a default greeting.

5.2 Control plane

Figure 3 shows the control plane. The control plane, like the data plane, has components in both networks. Note the presence of the telephony server, which is used to control the actions of the PBX using the JTAPI interface. Also, note the database server, connected to the Internet, which stores control information. The telephony server uses JTAPI to register itself as a monitor on the voice mail number (currently 1 607 254 5510). It is then automatically informed when a call arrives for this number. Specifically, the telephony server registers a Java routine with JTAPI. On an incoming call, the routine is called with an argument set to the ID of the incoming caller. The routine can also acquire touch-tone digits as they arrive. When the telephony server has all the digits of the extension, it sends a message to the gateway with this information. The gateway uses this to look up an Access database (using JDBC) to determine the name of the greeting file for that user, and the corresponding email address. If the user is found, the gateway plays the greeting file on the serial line, and this greeting is forwarded by the PBX to the user. If no user is found, an error message is played to the caller. Since the lookup on the database can take a fairly long time (5-10 seconds), we added an additional handshake between the telephony server and the gateway. When the telephony server contacts the gateway, it puts the call on hold, so that the caller hears a ringing sound. The gateway

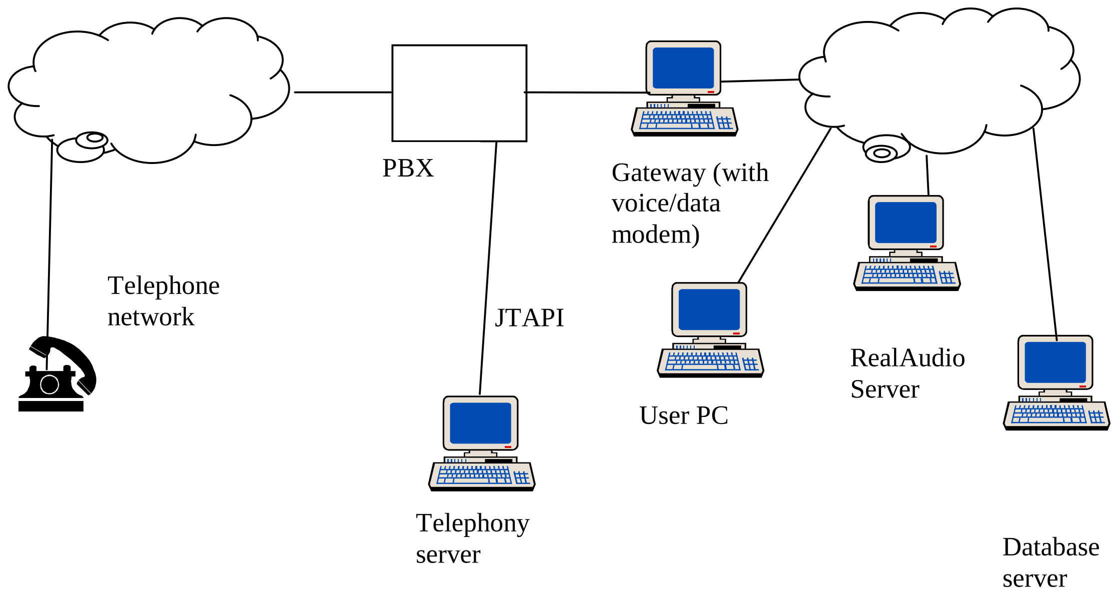

Figure 3: Control plane in the Internet voice mail application

starts looking up the database when it receives a handshake message from the telephony server. When the lookup completes, it sends a handshake message back to the telephony server. At this point, the call can be forwarded to the gateway's modem. When the telephony server receives the call disconnection message from the PBX, it sends a disconnect message to the gateway. The gateway then closes the file containing the digitized message and sends it, using SMTP, to the mail server.

The control actions required to store a greeting file are similar. The user dials a special access number, followed by the extension and an access code. When the telephony server learns of this call, it places the call on hold and sends a handshake message containing the extension number and the access code to the gateway. The gateway looks up the database to check if the access code is valid. If it is, it sends a handshake back to the telephony server. At this point, the call is forwarded to the gateway, which stores the incoming message as a greeting. When the call terminates, the greeting file is closed, and its name is registered with the database.

Note that we have partitioned the control functionality between the telephony server and the gateway. This is because we wanted to cleanly separate all interactions with the PBX on a single machine. It also allowed us to debug the Internet data transfer portion of the application separately from the telephony portion. We can easily combine all the control functionality in the same workstation, should that prove necessary.

Note also that in our implementation, we have a single connection between the Internet and the telephone network. This is sufficient for our experiment. In a larger deployment, it would be necessary to replace this single line with a set of lines in a trunk. Telephony cards that provide this capability are commercially available from suppliers such as Dialogic Inc. and we are currently working on integrating such a card into the existing infrastructure.

6. Discussion and previous work

The main contribution of this work is to describe the architecture of multiplanar applications and multimode networks. We believe that such applications and networks will become very attractive in the near future since they allow current providers to leverage their existing user bases in providing value-added services. Since users do not need to learn a new interface, and new services can be introduced nearly transparently to users, these applications allow network providers to evolve their technology rather painlessly.

Telephony gateways are well known in the literature, and have been used for many years in call-centers. However, these descriptions, to our knowledge, do not present the overall architecture of multiplanar applications that we have presented here. Nor do they extend to interconnection of all internetworking technologies, as we have done in our work.

The use of gateways to connect disparate datalink layers is well-known in the literature. Indeed, the Internet is precisely a way to interconnect networks that use different datalink layer technologies, and different administrative domains. In our work, we have extended this paradigm to include different internetworking technologies. We believe that this next step can lead to a variety of useful applications.

Similar motivations can also be found in the work by Katz and Brewer in their description of an *overlay* network [KB 96]. In their work, they describe a network architecture where a wireless terminal receives seamless connectivity as it moves from one wireless network to another, by appropriate handoffs between network gateways. Our inter-planar gateways are similar in function to what they describe.

8. Conclusions

In this paper, we have presented the evolutionary trends that lead to multiplanar applications and multimode networks. We describe characteristics of such applications and networks, and show why they have an advantage over uniplanar applications. We have presented several non-trivial examples of multiplanar applications, and have also described the design and implementation of one such application in detail. This application, while demonstrable, is not particularly stable, nor is it capable of supporting more than a handful of users. In current work, we are making the technology more robust, adding more lines to the gateway, and are tuning performance. Nevertheless, even as it stands, the application serves to demonstrate the intrinsic capabilities of multiplanar applications.

9. Acknowledgments

I would like to gratefully acknowledge financial and equipment support from Lucent Technologies. The Internet voice mail system described in Section 5 was done as part of a Masters of Engineering project by Kwan Hong Lee and Jonathan Lin at Cornell University. My thanks to them.

9. References

[KB 96] R. H. Katz, E. A. Brewer, "The Case for Wireless Overlay Networks," Proceedings 1996 SPIE Conference on Multimedia and Networking, MMCM '96, San Jose, CA, (January 1996).