52

# Directory Services for Internet Telephony: Creating a Spanning Layer over the Internet and Telephone Networks

by
Andrew Lester Sears

BS in Electrical Engineering
University of Missouri-Rolla, 1995

Submitted to the Department of Electrical Engineering and Computer Science
and the Technology and Policy Program
in Partial Fulfillment of the Requirements for the Degrees of

Master of Science Electrical Engineering and Computer Science
and
Master of Science in Technology and Policy

at the
Massachusetts Institute of Technology
September, 1997

Signature of Author ______________________

Certified by ______________________
Lee McKnight
Lecturer, Technology and Policy Program
Thesis Supervisor

Certified by ______________________
David Clark
Senior Research Scientist, Laboratory of Computer Science
Thesis Supervisor

Accepted by ______________________
Professor Arthur Smith
Chairman, Committee on Graduate Students

Accepted by ______________________
Richard de Neufville
Chairman, Technology and Policy Program

© 1997 Andrew Sears. All rights reserved.
The author hereby grants to MIT permission to reproduce and to distribute publicly this paper and electronic copies of this thesis document in whole or in part.



OCT 29 1997

**Directory Services for Internet Telephony:**
**Creating a Spanning Layer over the Internet and Telephone Networks**

by

Andrew Lester Sears

Submitted to the Department of Electrical Engineering and Computer Science, and the Technology and Policy Program on August 8, 1997 in Partial Fulfillment of the Requirements for the Degrees of Master of Science Electrical Engineering and Computer Science, and Master of Science in Technology and Policy


**Abstract**

The recent growth of Internet telephony and multimedia conferencing has created the need for directory standards to support these applications. This includes directories for locating users, real-time groups, hop-off gateways to interconnect with the Public Switched Telephone Network (PSTN) and other intermediate services. This thesis examines the issues involved in providing these directory services, proposes a design and models the scalability of that design.

The main challenge in designing the directory is that much of the information in the directory is dynamic. Because of its rapidly changing nature, traditional methods to provide scalability such as replication and caching would not be effective. The design proposed modularizes the dynamic and static portions of the directory to allow replication of static data while isolating dynamic data. The scalability of the dynamic dat is then achieved by partitioning the data across several servers.

The design proposed is implemented through the Lightweight Directory Access Protocol (LDAP) to the X.500 directory. Methods for providing directory services through DNS and HTTP are also presented. The design presented allows users to be located using an E-mail address or other personal identifier. It also allows users to create groups and provides mechanisms for creating persistent groups based on the Usenet structure. It provides a mechanism for gateways to the PSTN to register themselves, and for users to locate appropriate gateways to place calls from the Internet to the PSTN. It also suggests an automated "yellow pages" framework that other intermediate service providers (i.e. speech recognition libraries) can register for users to locate them.

This thesis also examines the policy and economic issues related to providing directory services for conferencing. One key issue is the privacy implications of these directory services, so methods are suggested to avoid violating individual privacy. Another issue is how directory services for multimedia conferencing will allow for the Internet to provide telephony services decoupled from the underlying network. This can be contrasted to the coupled nature of the PSTN in which the carrier has control of the services because they are provided by the switches controlled. This thesis examines these issues to consider the economic effects of providing these directory services.


**Thesis Supervisor:** David Clark, Senior Research Scientist, Laboratory of Computer Science
**Thesis Supervisor:** Lee McKnight, Lecturer, Technology and Policy Program

**Acknowledgments**

I would like to thank Mark Handley, Karen Sollins, Henning Schultzrinne, Frans Kaashoek, William Lie from Microsoft, Tony Genovese from Microsoft, Rich Pizarro from Netscape and Katie Roper from Four11 who offered comments on early ideas. Special thanks goes to David Clark who has provided continual feedback throughout the process of this work and Lee McKnight who has offered continual feedback on the policy implications of directory services.

This thesis has been supported the Department of Defense NDSEG program,
Contract # N0014-95-1-0265.

## Table of Contents

## Table of Figures

## 1. Introduction

This thesis examines the issues involved in providing directory services for Internet conferencing. The hypothesis considered is that existing Internet directories do not meet the needs of conferencing applications, and that directory structures based on the Lightweight Directory Access Protocol (LDAP) could meet these needs. The thesis is demonstrated by outlining the requirements for these directory structures, proposing a design for these directories in LDAP and modeling the scalability of that design. One of the key problems in designing these directories is that the data stored in them is dynamic, so traditional methods to achieve scalability such as replication and caching are not effective. The thesis demonstrates a design that modularizes the dynamic data so it can be partitioned across several servers.

The primary focus of this thesis is to examine the problem of providing a directory services for conferencing applications. This thesis presents a design of a directory service for conferencing based on LDAP. These directory services include such features as the real-time location of a user, the real-time location of members in a user group and the location of gateways for interconnecting the Internet with the PSTN. The secondary focus of this thesis is to examine some of the issues related to the convergence of the Internet with the PSTN. The thesis begins by presenting an overview of existing multimedia conferencing standards in order to explain the context of directory services for conferencing among other standards. The thesis then examines some of the different types of directory services for conferencing and reviews previous work in this area. The thesis then examines potential directory access protocols and explains why LDAP was chosen in the design. The thesis then presents the design of a schema for the directory and the scalability of the schema is modeled. Finally the policy and economic implications of directory services for conferencing are presented, and the thesis is concluded by examining future areas of study.

## 2. Overview of Multimedia Conferencing Standards

This section provides some background on existing conferencing standards to show how the conferencing directories suggested in this paper might fit in the existing standards framework. The first section explains the relevant ITU conferencing standards focusing on the H.323 and T.120 series. The next section examines the IETF conferencing standards. The third section attempts to bridge the two groups of standards and shows how there is a gap for conferencing directories.

### *2.1 ITU Conferencing Standards*

The key standards from the International Telecommunication Union (ITU) for Internet conferencing are the T.120 and H.323 series of standards [T.120] [H.323]. The focus of the T.120 standard is on data protocols for multimedia conferencing, and it specifies the relation of the standards in the T.120 series, which come from Study Group 8 of the ITU. These standards also provide support for multipoint communications using a centralized mechanism called a multipoint conferencing service (MCS). The focus of the H.323 standard is on video conferencing on local area networks (LANs) without quality of service (QoS) guarantee, which includes most LANs on the Internet. H.323 specifies the relation of standards in the H.323 series, which comes from Study Group 15 of the ITU. H.323 is basically an Internet "patch" of the H.320 series of standards which specify videoconferencing over the PSTN. The T.120 and H.323 standards are complementary although it appears that there is an overlap between some of the call control functions. A list of standards in both series is as follows:

### 2.1.1 H.323

Figure 1 shows how the components of the H.323 standard fit together. An explaination of the components are as follows:

1. **Video Codec**: Specifies the encoding method for video. H.261, H.263. Optional
2. **Audio Codec**: Specifies the encoding method for audio. G.711 (PCM), G.722 (7kHz speech at 48 and 56 kbps), G.723.1 (3kHz audio at 6.3 or 5.4 kbps), G.728 (3kHz audio at 16 kbps, G.729 (3 kHz speech at 8 kbps)
3. **System Control**: Specifies Signaling for call control, capability exchange, commands and indications and messages to open and describe logical channels. Specified in H.245, "Control Protocol for Multimedia Communications."
4. **Formatting:** Formats video, audio, data and control streams for output into the network interface. It also provides the logical framing, sequence numbering, error detection and correction for each media. This includes the use of RTP/RTCP. Specified in H.225, "Media Stream Packetization and Synchronization." All the standards in H.323 and T.120 can be considered to be "above" H.225 because they must all be formatted according to H.225 before going to the transmission protocol.

The following diagram shows how the ITU relates these protocols:

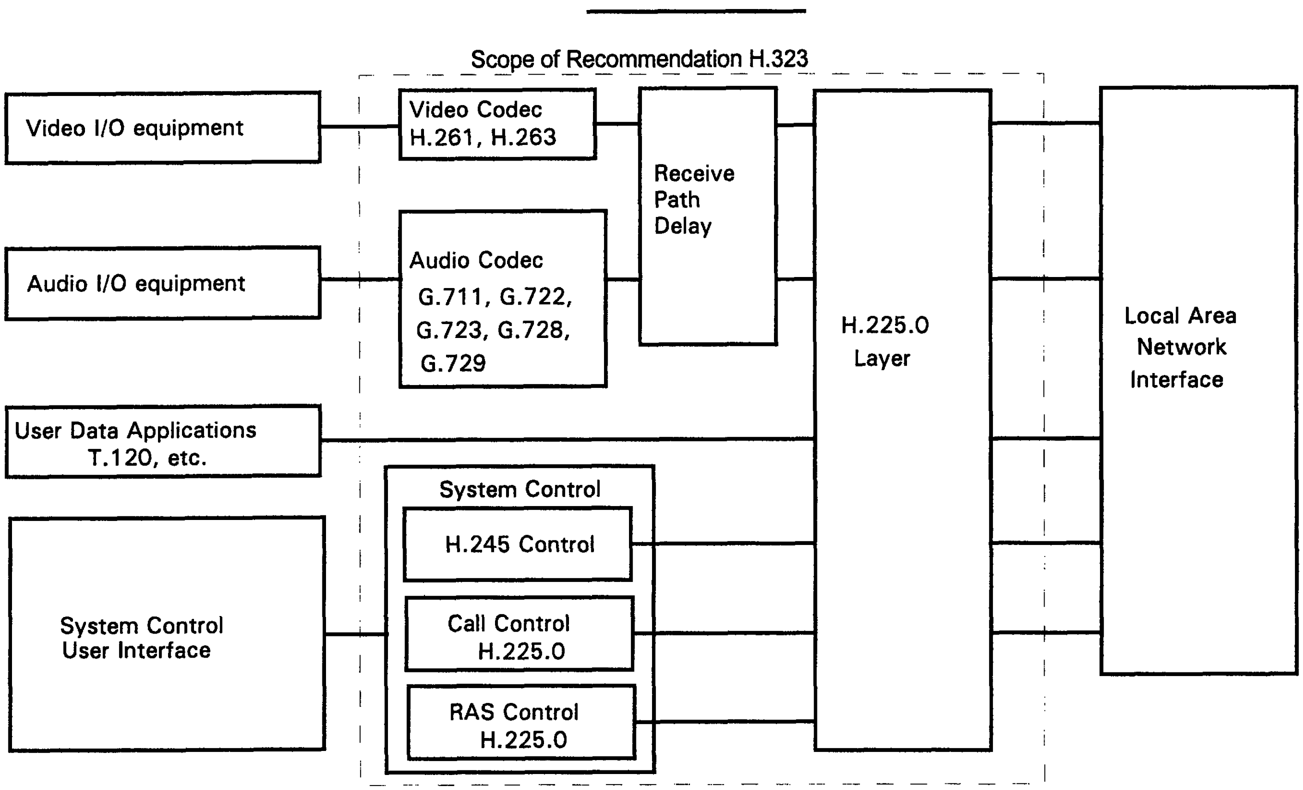


**Figure 1. Components of the H.323 Standard [H.323]**

### 2.1.2 T.120

Figure 2 shows how different protocols in the T.120 series relate. The T.120 series is complementary to the H.323 series and provides application level protocols and advanced conference control.

1. **Application protocols.** These standards provide application-specific standards that can be used for multipoint conferencing.
   a. File Transfer: T.127 specifies the standard for transferring a file in a multipoint conference.
   b. Shared Whiteboard: T.126 specifies the standard for sharing a whiteboard, which can be marked up by users in a multipoint conference
   c. Application Sharing: Microsoft has promised to promote a standard for sharing of applications that allows for the transmission of logical objects. (T.12?)

d. Generic Application Template: T.121 allows application writers to write their own application protocol for a unique application in conformance with the T.120 specifications.

2. Conference Control: Provides services for setting up and managing a multipoint conference, called Generic Conference Control and specified in T.124
3. Multipoint Communications Service (MCS): Provides a centralized mechanism for relaying a multipoint conference. This can be seen as an alternative to multicasting, which is not possible in the PSTN. The definition of the service provided by an MCS is specified in T.122, and the protocol of an MCS is specified in T.125.
4. OSI Transport Stack Profile: This specifies an OSI transport service interface to the MCS layer. This interface would need to be used for multipoint conferencing within the PSTN.

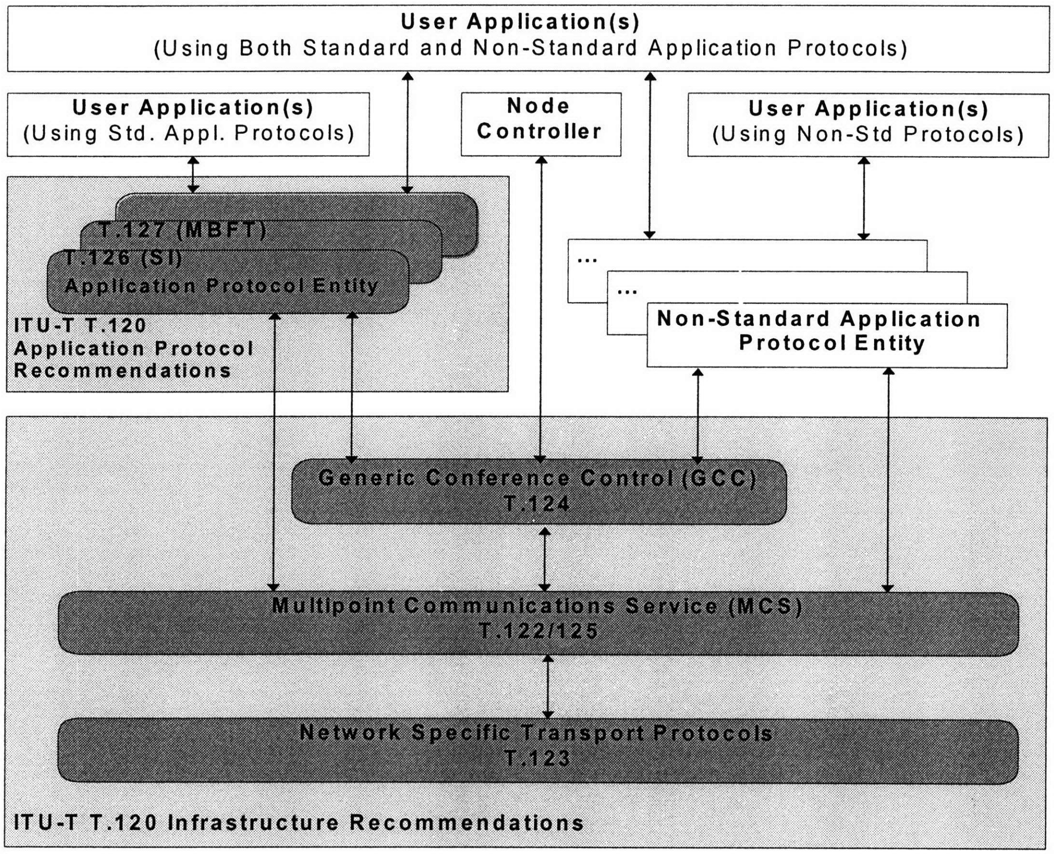


**Figure 2. Components of the T.120 Standard**

## 2.2 IETF Conferencing Standards

The IETF's focus multimedia conferencing has been primarily on standards to enable the use of multicast for multimedia conferencing. The standards are based on a framework called the Lightweight Session Architecture [FL95], which follows the flexible philosophy of the IETF called ALF [CL90]. The working group in charge of the higher level standards specifically for multicast conferencing is called the Multiparty Multimedia Session Control (MMUSIC) group. To have a full understanding of standards from MMUSIC, it is necessary to understand the underlying standards that they rely on from other IETF groups. A list of standards includes:

1. Internet Group Management Protocol (IGMP) [FE97]. Provides management functions for multicast routers, such as a router joining a multicast tree.

2. Resource Reservation protocol (RSVP) [Br97]. Provides the setup protocol to reserve resources in a router, which can be used to provide reliable QoS from routers.
3. Real Time Protocol and Real Time Control Protocol (RTP & RTCP) [Schu96a]. RTP is a short header which is added to UDP packets for carrying real time information such as audio or video. It includes time-stamp information for synchronizing audio and video, a sequence number for detecting lost packets and identifiers. RTCP provides control information such as feedback on the quality of data transmission. RTCP can also send Source Description (SDES) packets with information on members in a multicast group, which can be used to maintain a real-time list of membership in a conference.

**MMUSIC Standards**

1. Session Announcement Protocol (SAP) [Ha96c]. This specifies how users send a multicast announcements to a well known channel with information about multicast conferences. This provides a similar function as the Conference Query and Conference Create services in T.124. The difference is that T.124 requires a centralized MCU, while SAP provides some functions useful specifically for multicasting (such as zoning). T.124 currently does not provide multicast support, but if it did, then SAP could be combined with T.124. The conferencing directory standard described in this paper could be used to serve as a cache for SAP.
2. Session Description Protocol (SDP) [Ha96b]. Specifies the format of announcements sent using SAP. This overlaps with T.124 messages for Conference Query and Conference Create as described above.
3. Simple Conference Invitation Protocol and Simple Invitation Protocol (SCIP/SIP) [Ha96d] [Schu96c]. Both of these protocols specify the information to be exchanged to invite someone to a conference. These are competing protocols that will be combined. SIP is based on the format

of SDP, while SCIP provides more information than SIP and uses a slightly different format. This standard is in competition with some of the functions in H.245 (such as exchange of terminal capabilities), and would not be needed if H.245 becomes the dominant standard.

4. Simple Conference Control Protocol (SCCP) [Bo96]. SCCP specifies a control protocol for tightly coupled conferences. Functions include management of members in a conference, management of applications/media in a conference, floor control, and moderator control.

Figure 3 shows the layering of the various IETF conferencing standards. A standard which is above another standard indicates that it runs "on top of that standard."

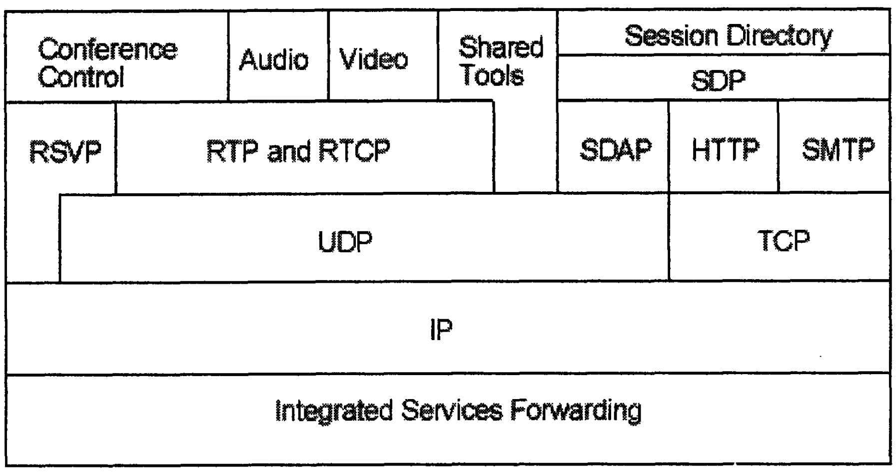


**Figure 3. IETF Conferencing Standards**

### *2.3 An Integrated Framework*

It is difficult to show exactly how the ITU and IETF standards compare directly to each other. This is largely because the ITU standards are monolithic, while the IETF standards are divided into the function they serve. Many of the conferencing standards of the IETF and the ITU are complementary, but a few overlap. Before overlapping standards can be "weeded out," it is first necessary to make some assumptions about what the world will look like at least in the short term for multimedia conferencing. Three assumptions include:

1. **Multimedia conferencing will be done both through the Internet and the PSTN.** This is largely because it may be some time before the Internet can provide QoS suitable for all users.
2. **Client-based, centralized MCU and multicast methods of multipoint conferences will be needed.** Client-based multipoint conferences (where the client sends multiple packets) are the easiest to implement, and are available now, but this is inefficient and will not scale well. MCU's will be needed because it will be some time before multicasting will be widely available on the Internet and is likely to never be available on the PSTN. Multicasting is needed because it is likely to provide increased efficiency over MCU's.
3. **H.323 seems to be the dominant standard in its domain.** With the support of Microsoft, Netscape and VocalTec among many others, it appears that something similar to H.323 will become the dominant standard. Exactly what this will mean remains uncertain because an industry group (VoIP) has been formed to work out the implementation details.

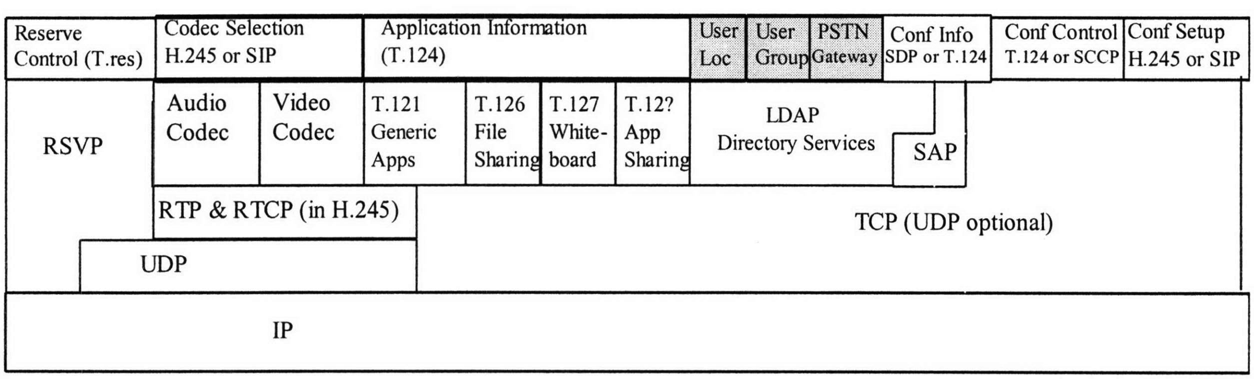


**Figure 4. Integrating IETF and ITU Standards**

The diagram in Figure 4 attempts to combine the two groups of standards and show their layering. Each higher layer is dependent on the layers below it. Adjacent protocols should not have any dependency between them; however, the conference setup protocol may be dependent on information to locate a conference if a directory is used. Most of the standards shown in the diagram were described in the previous section. The top layer provides control functions and the exchange of information. In this layer, the monolithic protocols H.245 and T.124 are broken down into the various functions to better show their relation to IETF standards. The various functions at this level are as follows:

1. Reservation Control. This would specify how a higher level conference control or setup protocol would use RSVP to reserve resources. This is currently under study to be included in the T.120 series.
2. Codec Selection. This specifies how information on codec capabilities are exchanged, and possibly the procedures for codec selection.
3. Application Information and Control. This specifies how information on application capabilities are exchanged, and possibly the procedures for activating an application.
4. User location, User groups and PSTN directories. These are the standards proposed in this thesis based on LDAP.
5. Conference Information. This provides information on conferences and their availability.
6. Conference Control. This is for all control functions not specified elsewhere.
7. Conference Setup. This specifies the information exchanged and the procedure for bringing someone into a conference.

Figure 4 also shows the relation of the directories that are being proposed in this paper. The directories proposed in this paper would run on LDAP, which would run on TCP or UDP over IP. As shown in the diagram, they would essentially be independent of other standards.

# 3. New Directories Needed for Conferencing

There are several different functions needed by conferencing directories. These functions include locating a user, finding groups of users and selecting a user in a group and locating gateways for connecting with the PSTN. This section will provide details on each of the functions of these directories, listing the steps involved with each.

## *3.1 User Location Directory*

A *user location directory* serves to provide the location (IP address) of a user as well other dynamic information such as whether they are logged in and available or not. *White pages* directories are related to user location directories. White pages directories contain relatively static information about a user such as their name or address, whereas user location directories contain dynamic information, such as where they are located and whether they are currently available. In the directory design proposed in this thesis, the white pages entry of a user either contains a subentry that contains the user's location or a pointer to the user location record.

If you want to call a friend on the Internet using a conferencing application, placing a call may seem easy, but there is a series of steps that works in the background that is much more complex. There is a series of steps that you must go through to complete your call which include:

1. Locating an identifier (uid) that uniquely identifies your friend. This might be an E-mail address, a phone number, an X.500 distinguished name, or domain name (identifying an individual rather than a machine i.e. andrew.sears.mit.edu) or any other naming system that might be available. This could be done manually by looking up the friend's name in a White Pages service like Four11, or simply by people giving each other their address.

2. Use the identifier to determine the host with a directory that knows that user's location. This step is needed because otherwise you would need to go to a centralized location to find that information, which would not scale well. The location of the directory may or may not be evident from the name. For example the identifier, foo@directory.mit.edu, might indicate that the machine directory.mit.edu contain that user's information. This would make the second step trivial. Another alternative is that the identifier indicates a domain of the directory, such as foo@mit.edu. Here you would ask a directory at mit.edu where its locator server is, and it would return something like directory.mit.edu (this is how mail is done). The third possibility is that the identifier is *location independent* in that it contains no information about where the users directory might be. In this case either there is a centralized location of this information, or some algorithm that can map identifiers into names of directory servers.
3. Contact the directory to find whether the user is logged in and where they are located. Here you might send a request the host directory.mit.edu, indicating the uid. It will then return the IP address, port and other information needed to contact the user. In the case of multicast conferences, this will be a multicast address.
4. Contacting the user. A call setup protocol then takes over and the call is maintained by standards in the T.120/H.323 series or the Lightweight Session Conferencing Architecture (LSCA) of MMUSIC.

One thing to note is that not all calls will go through these steps because in some cases, users may already contain the information provided in one of the steps. If you want to contact your friend and already know their IP address/port and that they are logged in, you could skip directly to Step 4, which is what many applications do. The previous three steps are needed to make calling more user friendly because most people are not likely want to prearrange all calls, and they may want to receive calls at different locations and use simple names like E-mail addresses to identify users. To provide these features, it is clear than additional standards for a user location directory is needed.

### *3.2 User Group Directory*

*User Group directories* can be thought of as a "virtual room" where people meet each other to start a conference. A large percentage of existing Internet conferencing calls are made using user groups because it allows people to meet new people on the Internet to talk. VocalTec and others have used only the directory functions of IRC (Internet Relay Chat) as a group directory. For example, in VocalTec's Internet Phone, a user might create a group called "sports" in which hundreds of people might join. Members in that group could then call each other in separate conferences to talk about sports.

The term *conference directories* refers to a "auditorium" where users could join a discussion on "sports" listen to it for a while, and then pull aside several members to set up a different conference. *Conference membership directories* maintain membership of a group after a conference has already been started. In a user group directory, all members may be talking in the same conference or none may be talking together or anywhere in between. However the term "conference" implies that everyone is involved in the same discussion. The key difference between conference and user group directories is that being a member of a group does not indicate that the user is involved in a discussion with members of that user group, whereas in a conference it does. In other words, the membership function of the user group is decoupled from who is participating in a discussion. Standards for conferencing directories are being provided both by the IETF (SDES packets in RTCP) and the ITU (T.124 Conference Roster). *Conference information directories* provide information on what conferences are available to a user. This information might include the conference title, topic of discussion, times, etc. This is analogous to the group information portion of the user group directory proposed in this thesis, with the same differences as described above for conference membership.

Two paradigms can be used to think of group directories: *IRC channels* and *Usenet groups*. IRC is currently used for text chatting on the Internet. In IRC channels any user can create a group, and there are many groups which are not very structured and do not scale well. Usenet is currently used as an electronic "bulletin board" where users can post messages to others on the Internet. In Usenet, users can navigate through groups by topic and subtopic, until arriving at the one they want. Navigating through groups might involve the following steps:

1. Locating a list of group topic areas.
2. Navigating through subtopics and receiving group information. This might include requesting information on globally replicated groups, or groups that are only stored locally on another server. Information on groups might include group description, location of the server which contains the member listing for the group, access restrictions, etc.
3. Retrieving information and location of members in a topic or subtopic. This might include providing many users' nicknames, an E-mail addresses or UID's, comments (i.e. "Want to talk baseball") and their location information. This is essentially the user location function with some additional information and it provides information on all users in a group.
4. Contacting the user. Just as with user location, a call setup protocol then takes over and the call is maintained by standards in the T.120/H.323 series or the Lightweight Session Conferencing Architecture (LSCA) of MMUSIC.

### *3.3 PSTN Gateway Directory*

PSTN gateway directories provide information needed to place calls from the Internet to the telephone network. The steps involved in using a PSTN gateway include the following:

1. Locating the telephone number
2. Requesting information on gateways that can complete the call for a given service. Given a phone number, this directory will produce a list of gateways that are close to the number being called that can complete the call. Information in the directories is needed to specify what type of gateway is being provided, voice, fax, pager, data, etc. Information might also be needed on the price of using the gateway.
3. Contacting the gateway. Just as with user location, a call setup protocol then takes over and the call is maintained by standards in the T.120/H.323 series or the Lightweight Session Conferencing Architecture (LSCA) of MMUSIC. This will provide negotiation of protocols to be used for the call.
4. Placing the call on the PSTN. This may be done using regular dialtone (DTMF) or using advanced features through SS7.

## 4. Review of Previous Work

Table 1 shows the various directories related to conferencing. The first three columns for user location, group membership and gateways are the focus of this thesis. The directories in these columns will be discussed in further detail in the following sections. The second three columns show other directories related to conferencing that are not covered in this thesis because they are being worked on elsewhere. The rows show the protocols in the IETF, the ITU, miscellaneous protocols and the proposed protocols in this paper. For each of these, the access protocol, which is used to transmit the information, is listed separately from the specification of how the attributes in the directory are stored.

**Functions of Conferencing Directories**

| | User Location | Group Members | Gateway | White Pages | Conference Members | Conference Information |
|---|---|---|---|---|---|---|
| **IETF: Access** | None | IRC | DNS | LDAP | RTCP | SAP |
| **IETF: Attributes** | None | IRC Members | tcp.int | White Pages | SDES | SDP |
| **ITU: Access** | None | none | none | 411, etc. | RTCP, T.124 | T.124 |
| **ITU: Attributes** | None | none | none | For further Research | SDES, T.124 Conf. Profile | T.124 Conf. Roster |
| **Misc: Access** | MUCS, ULS | none | none | HTTP | none | None |
| **Misc: Attributes** | SCIP, ULS | none | none | Four11, Whois, etc. | none | None |
| **Proposed: Access** | LDAP | LDAP | LDAP | LDAP | RTCP, T.124 | SAP, T.124 |
| **Proposed: Attributes** | Section 5 | Section 5 | Section 5 | White Pages | SDES, T.124 | SDES, T.124 |

**Table 1. Standards and Proposals for Conferencing Directories**

## *4.1 User Location*

There are currently over two dozen different Internet telephony applications, each of which has its own directory. None of these directories are interoperable with the others. Without a common directory standard, n different conferencing applications will require $n^2$ implementations to achieve interoperable directories. Existing conferencing directories also utilize flat structures, which limits their scalability. In the tests with the three most widely used applications, it was found that

Microsoft, Netscape's and Vocaltec's directory servers were often congested and dropped requests. A more robust, scaleable, and open design is needed than what these proprietary directories provide.

### 4.1.1 Multimedia Conferencing System (MUCS)

Probably the most thorough proposal to date for an end-to-end system for user location in the Internet is the Multimedia Conferencing System (MUCS) presented in [Schu96b]. It proposes to use E-mail addresses to uniquely identify individuals. These addresses would be used with the SRV records in the Domain Name System (DNS) to provide redirection to a designated conferencing directory server for that domain. The conferencing directory server would use the MUCS access protocol to the directory. It also suggests a record format called the Simple Conference Invitation Protocol (SCIP) [Schu96c]. One problem with the MUCS proposal is that it requires the creation of a new directory access protocol (namely, MUCS). This means that MUCS would only be able to provide directory services to user domains that have MUCS servers.

### 4.1.2 User Locator Service (ULS)

Microsoft has proposed a standard called User Locator Service (ULS) [Wi96] which functions in a similar manner as MUCS and suffers from the same problems. In particular, it creates an entirely new directory interface. Although Microsoft has published this interface, it has not made source code available for its implementation. As is often the case with other proprietary systems, this failure to release source code limits its development and adoption. By itself, the ULS proposal is also incomplete because it does not specify a set of attributes for conferencing records and does not indicate how it might be linked to a global namespace such as DNS or X.500.

### 4.2 User Groups: IRC

Internet Relay Chat (IRC) [Oi93] provides a mechanism for providing group listing and membership directory information. IRC groups are an application-specific directory for text-based conferencing. Some conferencing applications, such as Vocaltec's Internet Phone have adapted IRC to provide group information for conferencing. However, the design of IRC has severe scaling problems because user and group information is globally replicated between servers that are linked together. The result is that IRC has broken into many "islands" of IRC servers which are not interconnected. Even so, many IRC "islands" still have scaling problems, and VocalTec's directory has shown poor performance from tests in the past.

### 4.3 PSTN Gateway: TPC.INT

A very basic form of a PSTN gateway directory called "tpc.int" has been implemented in the DNS [Ma93]. Tpc.int will return the domain name of a gateway given a phone number. It does this by creating a domain in DNS called tpc.int. Phone numbers are stored in reverse order in the DNS. For example a gateway that can call phone numbers with the prefix 617-225 would have the following record in the DNS *.5.2.2.7.1.6.tpc.int. A user request to the DNS would look like, 4.3.2.1.5.2.2.7.1.6.tpc.int. The asterisks would serve as a wildcard, and would return the IP addresses of gateways that can serve that number.

The problem with tpc.int, is that it does not provide any information about that gateway (such as the price to use it), and it does not distinguish gateways based on the type of service they provide. Tpc.int is currently being used mainly as a directory for fax gateways, so users desiring voice gateways to the PSTN have no way of distinguishing voice gateways from fax gateways. Because of the many different types of gateways and the range of functionality supported by each, it is likely that an open gateway directory will need more functionality than can be provided through the DNS.

## 5. Directory Design: Selection of Directory Access Protocol

The first decision in designing conferencing directories is the selection of a directory access protocol. The *directory access protocol* specifies how information in the directory is accessed and transmitted, and controls things such as access rights, commands to access data and defining how lower level protocols will be used to transmit the data. This section details the functionality needed in an access protocol and lists the possible choices. The section concludes by explaining why LDAP was selected as the access protocol.

### *5.1 Functionality Needed in an Access Protocol*

In designing conferencing directories the first question to consider is what directory access protocol to use. An examination of existing conferencing applications was used to determine the needed functionality for an access protocol to access directory information. The functionality of existing conferencing directory access protocols is currently limited to reading, modifying and creating user location records, forming new groups as well as reading and joining groups. Very few applications have any type of access control or mechanisms to verify a user's identity. In addition, these

| | Ideal | LDAP | DNS | HTTP |
|---|---|---|---|---|
| Global Namespace | Yes | X.500 DN | Domain Name | URL |
| Read* | Yes | Yes | Yes | Yes |
| Write* | Yes | Yes | Soon | Yes |
| Search* | Yes | Yes | By Name | No |
| Type Specific Access* | Yes | Yes | Limited | No |
| Object Oriented | Yes | Yes | Limited | No |
| Replication* | Yes | Yes | Yes | Yes |
| Authentication* | Yes | Yes | Soon | Difficult |
| Read Access Control* | Yes | Yes | No | No |
| Dynamic Optimized | Yes | Soon | No | No |
| Dyn Server Avail | Yes | Yes | Unknown | No |
| Internet & PSTN | Yes | Yes | No | No |
| Protocol Change | No | No | Yes | Yes |

**Table 2. Desired Access Protocol Features**

applications have their own proprietary access protocols to local, flat directory structures, with no connection to a global namespace.

Rather than take the limited view of functionality provided by these proprietary access protocols, this thesis provides an examination of the functionality of more developed access protocols such as those used by LDAP, DNS and HTTP. By examining these protocols, a list of potential functionality was developed: a global namespace for distributed operation, read access, write access, searching, type specific access, authentication, caching, replication, access control (read/write), whether the protocol provided optimizations for dynamic records, whether a protocol change would be required, and object oriented design. The entries in Table 2 marked with a "*" are essential to a conferencing directory service, while others are simply advantages. This table will be explained in full detail in the following sections.

## *5.2 Choices for an Access Protocol*

### 5.2.1 Dynamic LDAP

The Lightweight Directory Access Protocol (LDAP) is a simplified access protocol to the X.500 directory [X.500, We92] for the Internet. X.500 is a hierarchical directory, and although it has similarities with the DNS it has a different naming convention. X.500 records are identified by their Distinguished Names (DN's), such as "CN=Andrew Sears, OU=EECS, O=MIT, C=US." The last three parts of the name indicate different levels of a server hierarchy as shown in Figure 5, where each represents a server that is responsible for identifying servers and entries below it. For example, the diagram shows a user in

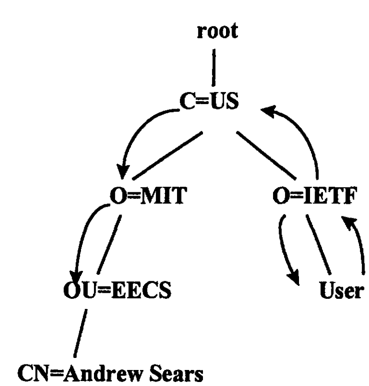

**Figure 5. An LDAP Request**

the O=IETF contacting the IETF server requesting the record with the above DN, and that server then contacts or redirects to the C=US, then the O=MIT server, then the OU=EECS server, which finally returns the record. This example shows how the system works without replication or caching, but normally higher levels of the hierarchy would be widely replicated and the location of frequently accessed servers would be cached to improve performance. LDAP specifies the interface to the X.500 directory, and provides a simplified protocol to communicate with servers.

Recently a standard was proposed to modify LDAP in order to handle dynamic records [Ya96]. It offers the greatest functionality and flexibility of the available access protocols, and LDAP is receiving widespread commercial support. Conferencing directories could be implemented in LDAP without changing the protocol using existing servers; however, existing servers have not been optimized to handle dynamic data and may perform poorly. There appears to be support for dynamic LDAP, considering that Microsoft has recently released their LDAP server optimized for dynamic data. Another significant advantage of LDAP is that it is the de facto standard for the directory for Internet White Pages, which could allow for seamless integration of user location and White Pages information.

After reviewing the above access protocols in detail, it was concluded that LDAP provides the best choice for the primary directory access protocol for conferencing directories. This is not to say that LDAP should be the only protocol used to access these directories. It is possible to provide gateways to other access protocols, and to use multiple access protocols linked to the same back-end database on a machine. For example, both Microsoft and Netscape provide a httpd server which shares the same database as LDAP and provides records using HTTP in addition to LDAP. This type of integration could also be done for parts of the DNS if it was desirable. This would allow the same objects to either be named with DNS or X.500 naming conventions, with the back-end database handling concurrency issues. The design specifies a conferencing directory using LDAP in the X.500 namespace. It could also be adapted for DNS, however with less functionality than LDAP.

### 5.2.2 Dynamic DNS

One possible directory access protocol is the Domain Name System [Mo86, Mo87]. Changes to the DNS have also been suggested which would allow for dynamic updates to the directory [Vi96]. These changes actually are designed to allow users to have write access to their DNS entries, rather than to optimize the directory for dynamic data as in Dynamic LDAP. The main goal of these changes is to allow laptop computers to maintain a correct DNS entry when they change their IP address as they are relocated. It might be possible to use the DNS as a general purpose dynamic directory, without changes to the specification. A major disadvantage of DNS is that it provides limited ability to specify attribute-value pairs, and does not provide the extensibility of LDAP.

DNS has the advantage of having many servers already deployed, but this advantage is not so significant because those servers would require a software upgrade before they could be used for dynamic data. The main problem with DNS is that it was never designed to serve as a general purpose directory as was LDAP/X.500, and there are currently no plans to use DNS to store White Pages information as there is with LDAP. Those modifying the protocol and developing new servers do not yet seem to have the intention of optimizing the protocol and server for dynamic data, as is being done with LDAP.

### 5.2.3 HTTP

HTTP [Fi96] is another possibility for an access protocol to conferencing directories. Its main advantage is its ubiquity, which would allow directories to be easily viewed by anyone with a browser. The main disadvantage of HTTP is that it was designed as a document access protocol rather than a directory access protocol. It does not include features such as type-specific access to attribute information. While it provides authentication for reads, authentication for writes requires the setup of individual user accounts on each server which makes implementation difficult. Because HTTP is not designed as a directory access protocol and lacks much of the functionality needed, it cannot be considered an appropriate choice for a conferencing directory access protocol, although it could provide read access to the directory.

### 5.2.4 Design a New Interface

Another possibility is that no directory access protocol is suited for conferencing directories, and a new directory interface should be designed for conferencing directories. This is the approach taken by the MUCS and ULS proposals. The main criticism of this approach is that it is quite a significant task to design a new directory access protocol, and there is no reason to "reinvent the wheel" if an existing protocol is appropriate.

### 5.2.5 RAS in H.245

Another possible access protocol is to use RAS signalling in H.245. RAS can be thought of as a channel that is used for control functions in a conference. RAS signalling is a function provided in the H.323 model for registration and admission of clients through a gatekeeper. This gatekeeper is a centralized mechanism to provide server support and control in conferencing applications. There is still widespread debate as to whether Internet conferencing application will need all the functions of the gatekeeper, and whether it will be widely deployed. One of the main problems of the gatekeeper seems to be its centralized design makes it complicated to implement. One problem with using RAS as the access protocol is that it is tied to the gatekeeper, which seems to have an uncertain future.

The second problem with the using RAS signalling as an access protocol is that it essentially requires building an entirely new data model and mechanisms. Until recently, RAS defined little more than a channel to send messages. There are already well developed data models and access mechanisms in LDAP which are likely to be able to meet the directory needs of conferencing applications, but RAS would require recreating data models. Because of this the argument against using RAS is similar to the argument against creating a new directory protocol: until it is proven that no other solution already exist, then building on RAS might be considered "reinventing the wheel." If some functionality is needed that cannot be met by existing technologies, then RAS may play a role with directory services and call management agents in the future. In the mean time, there are pressing needs for directories that need to be met soon, and another directory mechanism could provide parts of the solution. It remains to be seen what role RAS will play, but the proposal in this paper could be seen as either a complement or an alternative to providing directories and call management through RAS.

## 6. Directory Design: Selection of Organizational Schema

### *6.1 Overview of Proposed Organizational Design*

After an access protocol has been selected, the next step in the design of a conferencing directory system is the schema, which is the organization of the information in the directory. This involves consideration of what information will be available, who will maintain the servers for the directories, who will have update rights to the data, how different directories might interconnect, and how records can be replicated. While the proposed organizational design is specific to LDAP/X.500, it could also be adapted to other directories, such as DNS. Figure 1 shows the organizational design that is proposed.

It is assumed that user and group information will be stored under each organization in the X.500 tree, which is the directory structure for LDAP. This allows individual organizations to maintain their own directories. The records under the organization providing the conferencing directory can be logically divided into user records, group records and gateway information. The replicated section of the directory will include information for group information, pointers to users' user location record, and PSTN gateway information. The primary copy of all groups will be stored under an organization responsible for maintaining each particular group, but some of their information will be replicated. Since X.500 allows for multiple naming of records, it is also proposed that a new branch of the X.500 directory be developed under Internet information specified in [Jo94] to provide a secondary naming structure for these global groups, to aid navigation (explained in Section 5.2 and 5.3). The full explanation for this directory structure is in the following sections.

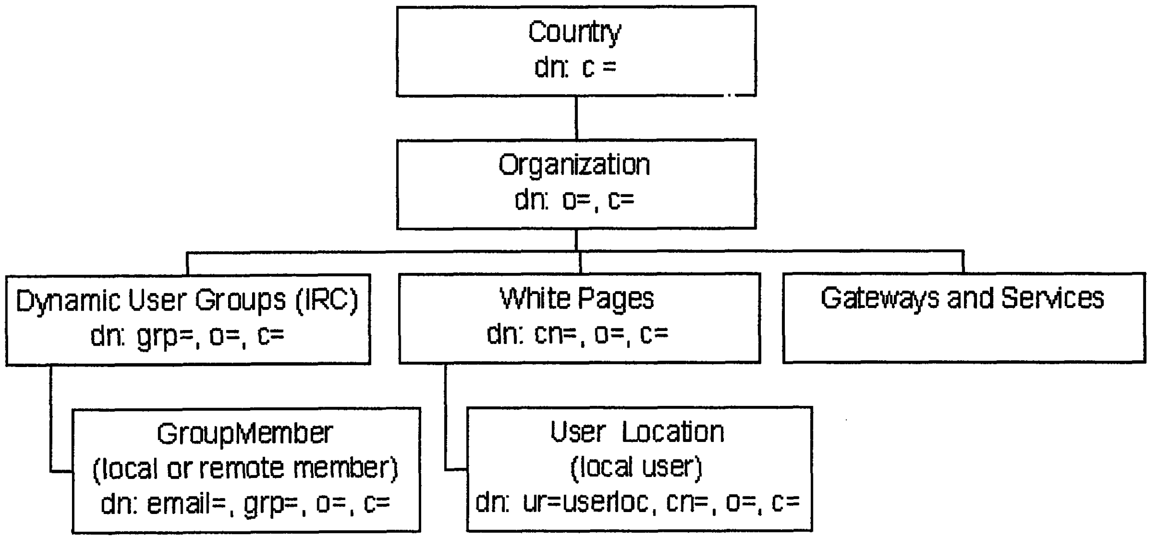


**Figure 6. Information Stored Locally in the Proposed Organizational Schema**

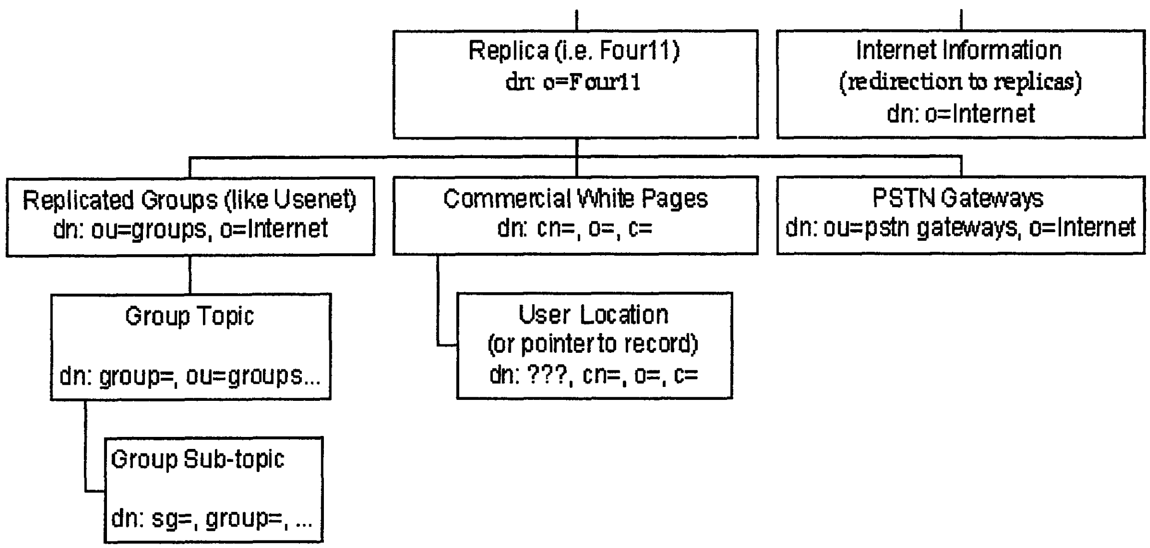


**Figure 7. Replicated Information in the Proposed Organizational Schema**

### *6.2 User Location*

In the design user location records would be stored as another record under each user's white pages entry or at a separate location. It is assumed that the white pages work will define some type of extensible record type under the white pages, which is termed "user record" or "ur." The DN of for the user location record will use the same DN as the user's white pages record adding the extension "ur=user location." For example the DN for John Doe's user locator record at MIT would be "UR=user location, CN=John Doe, O=MIT, C=US." It is likely that that some organizations might restrict access to this information. Some organizations may wish to have their user locator information server provided by their conferencing software provider or a third party (such as Four11). Here the DN might be "UR=user location, CN=John Doe 721, C=US." Currently all user location is done by application provider or third party, but it is likely that many organizations may choose to provide their own user location servers along with their white pages servers. The advantages and disadvantages of doing user location internal to an organization as opposed to by a software provider or third party is likely an issue that will have to be played out in the market. The design allows for whatever combination of possibilities that may result.

## *6.3 Group Directories*

### 6.3.1 Group Information and Group Member Records

Figure 6 shows that group directories actually consist of two records. *Group information records* store a group's description, access rights, location and other information. *Group member records* are very dynamic because they provide the real-time membership in a group using individual records for each user in the group, and include information such as the user's location (IP and Port), a nickname, a user's comments and DN of the user's White Pages entry. Group information records may be widely replicated, but group member records can only be replicated among a small group of servers if at all to maintain scalability. In the design, it is assumed that only the server storing the primary copy of the group information record will store group member records, unless another copy is desired for redundancy.

In considering group information records, the design makes a distinction between persistent and dynamic group information. This thesis defines *persistent groups* as those that users cannot add or delete without supervisory approval, so they do not change often. *Dynamic groups* are those that are created and deleted by users and change with relatively high frequency similar to IRC channels. The group information of persistent groups can be replicated and have multiple names which makes them "global" groups like Usenet. The group information of dynamic groups is in general not replicated, although it may be replicated within an organization or small groups of organizations. The primary reason for this structure is to improve scalability as will be explained is Section 6.

Our design of persistent groups is based largely on the Usenet model. It is expected that there will be primarily two ways to browse through group hierarchies: by organization and by topic. This corresponds to the two main branches "C=, O=" and "O=Internet, OU=Groups." A group identified by organization might have the DN "C=US, O=Microsoft, Group=My user group" while a group located by topic, might have the DN "O=Internet, Group=Sports, SG=Basketball." Browsing by topic allows users to see all subtopics in a group regardless of where the group is actually stored. The basketball group may actually be stored on a Netscape server and have a second name "C=US, O=Netscape, Group=Basketball." All groups, whether persistent or dynamic may be browsed by organization, but only persistent groups may be browsed by topic. This is because browsing by topic requires centralized information which translates the topic to the location of the server that maintains that group. This will only scale well if groups do not change often so their information can be widely replicated.

### 6.3.2 Usenet Groups as Conferencing Groups

This thesis proposes that existing Usenet groups and categories form the basis of the first conferencing groups, although future groups may not be limited to the Usenet structure. To do this, the topic area moderators under Usenet would need to be given passwords to modify their group information under the "O=Internet, OU= Groups" branch of the X.500 directory. For example, the moderator of the "alt.sports" newsgroup, would be given the password with access rights to modify and create new groups under the DN, "O=Internet, OU=Groups, Group=alt, Subgroup=sports"

While this answers the question of who will moderate a group, it does not necessarily solve the problem of who will provide the server to store the group. Providing the primary copy of a Usenet group involves relatively low overhead because the postings are only replicated at intervals, and can be "fanned" out using replication hierarchies. Providing the primary copy of a conferencing group involves maintaining real-time group member records for that group, which involves a higher overhead than maintaining a newsgroup. As explained in Section 6, the current cost might range from $.10 to $.40 per user, and are likely to decrease rapidly. However, this high of a cost is likely to detour many institutions such as universities from maintaining the servers. This leaves the application providers or a third party to maintain the servers.

It is possible that application providers will try to have their own closed groups that are not available to other applications. One way would be to have some type of access restriction on their servers through LDAP, that only their software can resolve. Another way would be for the application providers to not cooperate with the moderators of Usenet groups or to abuse their own power if they serve as moderators. It is likely that many application providers maintaining group servers may want to moderate the groups they maintain, which should be their right. In order to do this in an open manner it would require that they encourage browsing by topic rather than browsing by organization. Otherwise, Usenet conferencing groups would "balkanize" into Microsoft groups, Netscape groups and other application provider groups.

### 6.4 PSTN Directories

Another directory function that is needed by conferencing applications is one for locating gateways to interconnect with the public switched telephone network (PSTN). Given a phone number and the gateway service needed (fax, paging, voice, etc), this directory will provide information about gateways that can complete the call. This would allow calls to be placed from the Internet to the PSTN transparent to the user. This thesis proposes that gateway providers make the information they want available through their own LDAP server, making it open for replication.

### 6.5 Other Issues

#### 6.5.1 Gateways for SAP and T.124 Information

Both SAP and T.124 provide conference information that could be represented as group information in LDAP. Conferences can be considered a special type of group in which all members of the group must be in a conference. A gateway could provide conference information through LDAP (using session announcements in SAP and Conference Query in T.124). The same could also be done for conference membership (using SDES packets in RTCP and conference roster in T.124).

One reason this might be done would be to provide a common interface (LDAP) to all conferencing information. Another reason this might be useful would be to use the directory structure to organize conferences in ways not possible in SAP or T.124. One specific example of this is that SAP currently lacks scalability because it multicasts all announcement on one address.

The obvious solution to this problem is to set up a hierarchy of multicast announcement "channels." This hierarchy information could also be multicast through SAP, but it would probably be a waste of network resources since the hierarchy information is not likely to change often, making a server-based approach more appropriate.

The alternative solution would be to use LDAP to store the hierarchy, and either SAP or an LDAP gateway to SAP could be used to access announcements. This would require a new set of attribute-value pairs to match the SAP format. Another problem with SAP is that announcements are only received periodically, so it may take some time before a user receives all the announcements. The proposed solution to this problem is to have a server cache SAP announcements, although the proposal does not yet specify what access protocol will be used to get announcements from the server. This thesis proposes that LDAP be as used the access protocol for the server cache.

### 6.5.2 Fault Tolerance and Redundancy

From an examination of existing conferencing directories, the dynamic information in user, group, and group membership directory entries is not currently replicated. Future systems may choose to be more robust by providing redundancy. There are two ways that redundancy could be built into the system: by replicating dynamic data and by switching servers. The first is that replicas are kept of the dynamic data, so when a server goes down, the replica can provide the data. In a static environment storing replicas has a low cost, which may only involve some disk space and occasional network exchanges. For dynamic data, storing replicas may result in considerable CPU overhead for both the primary server and the replica to maintain consistency. If the data is dynamic enough, a single level of redundancy could double the number of servers needed.

The second way that redundancy could be built into the system would be for the clients to switch servers. This might use the "heartbeat" mechanism in Dynamic LDAP, where if a client does not receive response to their "heartbeat" message from the server they could contact the parent server with notification that the server was down. The parent server could then redirect all requests to the new server, which could request lost information from clients. In this method, information would be unavailable for a maximum up to the period of the refresh request including the response timeout. This is because after a server failure, when the client does not receive a response to its refresh request, it will send its information to another server and request updates from that server. This second method has the advantage of not requiring additional servers, but it would require a change to the LDAP protocol.

### 6.5.3 A Framework for Other Service Directories

A directory to locate gateways to the PSTN is only one example of a more general problem of locating services or objects. Services that might be used with telephony include, voice mail, speech recognition, text to speech, etc. The Open Distributed Processing model uses the concept of a "trader" [ITU-TS] which provides mechanisms for announcing and discovering objects. One possibility for future work in this area would be to attempt to use the "trading" architecture as a way to locate gateways and other services.

## *6.6 Comparison to Previous Work*

The method of *user location* in the proposed design differs from existing proposals such as MUCS and ULS in three ways. The first difference is that it does not specify a new directory access protocol, but instead utilizes an existing protocol, namely LDAP. Secondly, the proposed design separates the set of attributes for user location from the access protocol, to allow for the development of gateways to other access protocols, such as HTTP and DNS. Finally, the proposed design attempts to integrate user location with existing work on an Internet White Pages directory.

The method of handling *group information and membership* in the proposed design differs from existing proposals for *conferencing directories* in several ways. *Conferencing directories* maintain lists of conferences and members in conferences, and are usually not based on a directory server, but on a separate peer-to-peer protocol such as a conference control protocol. The primary difference is that while all members in a conference communicate with one another, this is not required for groups. The *Group directories* that the design implements are based on a directory server, and do not necessarily specify conference membership – although it can serve that function. A group can be described as a "room" where people can meet, but not everyone in a group need be in a conference with everyone else.

While the method of group information and membership in the proposed design can provide the functionality of IRC's groups, the two differ in several ways. One major difference is that the proposed design does not require the global replication of group information that limits IRC's scalability (instead, it only replicates pointers to the location of groups - this is explained in detail in section 5). IRC also couples together group directories and a specific mechanism for text-based conferencing. The design focuses only on the directory function and separates the directory access protocol from the set of attributes in a group listing record and group membership record. This separation makes groups application independent, so that record formats can be created for voice or video conferencing as well as for text-based conferencing. The proposed design offers more functionality in the access protocol and attribute listings as well as greater scalability compared to IRC and other group directories.

## 7. Results: Modeling Scalability

### *7.1 Determining Bottlenecks*

This thesis divides the costs incurred by the directory system into *processing cost, network cost, memory storage cost and disk storage cost.* To answer the question of which of these resources becomes the bottleneck, some parameters will be assumed to get an order of magnitude estimation. Assume that a server has 100,000 active, simultaneous users. Assume that each of these active users would read a record averaging 300 bytes from the directory once every 10 minutes. Assume that each user updates their record every 30 minutes with an average length of 100 bytes. Assume that the full record for each user is 500 bytes, and that these records are stored in RAM. Assume that each user is in one group, which has an average membership of 50 members, with five membership changes per minute per group requiring 50 byte messages. The load of the server is as follows:

Memory Needed: 50 Mb (users) + 5 Mb (groups)

Network and Server Throughput Needed: .44 Mbps (users) + .41 Mbps (groups)

Requests Needed from Processor: 222.22 req/sec (users) + 8,333 req/sec (groups)

There are several results that can be drawn from this example. The first result is that while memory usage may become high, it will not become a significant bottleneck, since RAM can be extended to 1 Gb and above on many existing systems. The second result is that network and server throughput is not likely to be a significant bottleneck, since networks much higher than this speed are common. From this example, it is easy to see that **the key bottleneck will be the number of**

**requests per second that can be handled by the server.** The most highly optimized servers now can only provide 100-1000 requests per second, showing 8,555 requests to be a significant bottleneck. Another key result is that **these requests are likely to be dominated by the cost of broadcasting changes to group membership.** This is because if there are an average of M members in a group, with a join/leave rate of R, then the number of requests for groups, N, scales according to $N = R*M^2$. Multicasting could help alleviate this problem significantly, and would bring the total number down to 389 requests per second, by scaling groups at N = R*M. This is because group membership information could be multicast to a group rather than requiring that the server send duplicate updates to many users. If multicasting is used, then the larger the average group size, the lower the server overhead for group membership.

### *7.2 Read/Update ratio and replication*

One of the key points of difference between conferencing directories and traditional directories is the update frequency of records. Different design choices may be needed when records become more dynamic, replication and caching become less effective ways to distribute load. Caching dynamic data will cause consistency problems, and at some point, the overhead for replication can become larger than the savings realized by spreading the load over several servers. Where this point is, will depend on the cost of replication, what level of consistency is demanded between replicas.

The previous section showed that processing costs are likely to be the dominant factor in scalability. Replication on N servers decreases the processing cost of reads by a factor of 1/N, but it increases the processing cost of replication. The total processing cost becomes as follows:

$$C_{Total} = C_{read} * \frac{N_{read}}{S} + C_{update} * N_{update} * N_{update} + C_{replicate} * K * N_{update} * S$$

Where, S = Number of replicate servers and K = average ratio number of updates for each replication operation, which is a function of the level of consistency demanded between replicas.

To find the optimal number of servers, the derivative of the cost with respect to S is set equal to zero to get the following:

$$\frac{dC_{total}}{dS} = \frac{-C_{read} * N_{read}}{S^2} + C_{replicate} * K * N_{update} = 0$$

$$S = \sqrt{C_{read} * N_{read} / (C_{replicate} * K * N_{update})}$$

The key result of this equation is that the optimal number of replicated servers is proportional to the square root of the ratio of the number of reads to updates. This equation can be used to find the point where replication causes more overhead than it saves by setting S = 1 and solving for the read/update ratio as follows:

$$\frac{N_{read}}{N_{update}} = \frac{C_{replicate} * K}{C_{read}}$$

If the read/update ratio increases beyond this point, then replication will become detrimental. It is important to note that replication serves both to distribute load and to provide redundancy, so it may still be desirable in some situations even if it increases overhead. Even before the point where replication becomes detrimental, the cost of replication increases significantly. With very static records, the main cost is for disk space since average processing cost of replication is almost zero.

Our design avoids the scalability problem with dynamic data, because it uses fine-grained partitioning to distribute load across several servers rather than replication. Each of these servers stores the dynamic information locally, with replicate servers only being used to provide redundancy if desired.

### *7.3 Scalability of Group Membership*

Section 7.1 showed that the cost of broadcasting group membership was a key bottleneck. The reason for this is that when a user joins a group, they must be given the full listing of all members of the group, and their group member record must be sent to all members of the group. If there are an average of M members in a group, with a join/leave rate of R, then the number of requests, N, scales according to $N = R*M^2$. The reason for the dominance of the broadcasting of group membership is that it scales proportionate to the square of M, while most other aspects of the directory scale linearly with the number of users. One way to resolve this problem is through the use of multicasting, which makes scaling linear with the number of members according to N=R*M. Another solution is that the membership of a group is only periodically refreshed at rate, T. This also allows linear scaling according to N=T*M. Even when this second method is used, server load can be quite high because T must be kept fairly low to keep records consistent.

### *7.4 Scalability of Location Independent Aliases*

Assuming that each indirection record were 20 bytes, one might expect tens of millions of records in the short term and hundreds of millions of records in the long term. This would require 100-1,000 megabytes of disk space. As long as these records were not changed more than once a year, then the average number of updates would be on the order of .1 to 10 per second, which could be handled by a centralized server. The number of reads for such a table could be quite large, so it would need to be widely replicated. Clearly, with enough money, the system could scale, so the question of scalability becomes a question of economics. One way to consider the economic question, would be to use the design to provide personal phone numbers which are location independent and never change. If it is assumed that there might be 500 million users globally, which each make an average of 1 call per hour at peak time, the generated load would be about 140,000 queries per second. If each server could handle 100-1000 queries per second, then 140-1,400 servers would be needed. This assumes that all calls will make use of aliases, which probably will not be the case. While this may seem like a large number of servers, this type of system could provide personal phone numbers across the world for $1.4-$14 million, assuming a server cost of $10,000. If such a system could be connected to the phone network, its benefits would clearly outweigh its costs.

### *7.5 Scalability of Group Information*

The key issue for group information is to consider to what extent can group information be replicated while still retaining scalability. The design achieves scalability by separating groups into two types: Static groups which are similar to Usenet and Dynamic groups which are similar to IRC which are created by users.

The scalability of static groups will be examined first. Let G = Total number of groups and R = Rate of change (creation/deletion) of group information. The rate R is under the control of the local server administrator, which will decide which group's information to add to their server as groups are announced in a similar way as Usenet (perhaps using the same mechanisms). Adding this information will only need to be done when the group description changes. The actual group membership for each group will be stored on the "home" server for the group, so an administrator "subscribing" to a group require no cost except for storage space. The number of updates will be proportional to G*R. Dynamic groups are similar to IRC, which has group information that changes more frequently. To examine scalability, let U = number of users in the system, and R be the rate of group creation/deletion for each user, since any user can create a group for dynamic groups. The rate of change of group information will be proportional to the number of users according to R*U. If dynamic groups are stored locally, then with U=100,000 users on a system assuming each user creates/delete R=.2 groups per hour. This will generate 5.55 requests per second, which is reasonable. If dynamic groups were globally replicated, then 500 million users would generate 27,775 requests per second which would overload all servers. Clearly, dynamic groups cannot be globally replicated in a scaleable way, although some could be replicated to a limited extent.

## 8. Policy and Economic Issues

### *8.1 Privacy Issues*

Privacy is one of the most important issues to consider in designing directory services for Internet telephony. User location directories could create significant privacy concerns for individuals as well as organizations because they indicate the location of a user at all times. Without any privacy protection, a computer system could be designed that could track the location of thousands if not millions of users simultaneously. This could allow others to know when individuals log in and out at work, when they log in at home, when and where they are logged in while traveling, etc.

There are several mechanisms that can be used to protect privacy. One popular mechanism is using call management agents (CMA) [Pe97]. Call management agents provide a "smart directory" function that could be used to select who can view an individual's user location information. The way CMA works is that users trying to locate another user must first authenticate themselves and can only view a user's record if they fit some specified criteria. This idea seems to work well in theory, but it may not work well in practice because it would require that everyone create complex lists of who can view their information. Not only would this be complicated, but it might unintentionally block important calls. CMA can also be used to block user location information based on organization, so that calls could only be placed within a company.

If Internet telephony follows user patterns in the telephone network, then it is unlikely that many people will have "closed" user location information, where only users in their "call list" can locate them. It is probably more likely that users will have an exclusive list, where specific users and groups are "blacklisted" and not allowed to locate the user. The problem with this is that this still leaves users open to "data mining" computers that may track their every move. Because these data mining computers are likely to follow easily identified patterns like many repeated queries on a user's information, it will be possible for the CMA or directory to identify data mining and block these queries. Some policy along these lines is likely to be needed in addition to user specified policies to protect privacy.

One major problem with all of these policies to protect privacy is that they require that a user or query can be reliably authenticated. Reliable authentication of users is still not available on the Internet, and is not likely to be widely available for quite some time. Because of this it is likely that privacy will become a major barrier to the widespread deployment of user location and CMA services. Even when authentication does become available, the reliability of authentication is likely to vary significantly from region to region. It is likely that future policies on user location may only allow authenticated users to get information and organizations/countries with poor authentication policies may be blocked as a group.

While reliable user authentication may reduce many privacy concerns, it may also create new ones. There may be some instances where a user wants to remain anonymous, and authentication might create new privacy concerns. For example a user might want to call a company to ask a question, and authentication provides the company with information about the user that can be used to build mailing lists and customer profiles that the user may not wish to disclose.

User authentication also brings up privacy concerns for user group directories. One can imagine a "data mining" computer that monitors the membership of all active public user groups. This computer could determine what groups any individual belonged to for any particular time. For example, a computer could monitor "sex" groups to determine who belonged to them at what time—information that could later be used to embarrass or even blackmail users. This problem could partially be addressed as it is in IRC by using aliases; however it might still be possible to link aliases to an IP address that might identify a user. This would be much more difficult to do on a mass scale than simply monitoring lists, so it is likely to be less of a privacy concern.

### *8.2 Reasons for a Lack of a Conferencing Directory Standard*

An important question to consider is why a conferencing directory standard has not yet been developed. One answer can be found by examining the history of conferencing standards. The main reason is that most of the conferencing work in the ITU focused on conferencing through the PSTN, such as the H.320 series of standards. These standard relied on the PSTN to complete calls, and did not need advanced directory services. The conferencing work in the IETF focused on conferencing through multicast on the Internet, and the work was through its MMUSIC group. Neither of these groups of standards was focused on developing standards for simple Internet conferencing. The H.323 standard was an attempt to try to bridge this gap, but it essentially took the H.320 standards and patched them to the Internet adding RTP/RTCP. This patch was less than complete, so an industry group called VoIP was formed to try to work out the details. Directory services were not included in H.323 because it was based on H.320, a centralized standard that does not need decentralized directory services.

Another answer for the lack of a standard can be determined by examining incentives to develop the standard. Telecommunications companies lacked the incentive to design conferencing directory standards because they lead to decentralized intelligence and may cause them to lose control and revenues from the network. This might be the reason why H.323 and T.120 standards do not yet have any standard for a decentralized mechanism for directory services.

The architects of the Internet may have also faced a different type of incentive problem. If Internet telephony were successful too quickly then one of three bad things might happen: the Internet might be overloaded and crash, the government would regulate the Internet or traditional telecommunications carriers might even view it as a threat and withdraw their support. The result was that the focus of the Internet Engineering Task Force (IETF) on multimedia conferencing was exclusively on the use of multimedia conferencing through multicast, which would be developed much more slowly and without threatening existing carriers. While this actually was an effective way to divide standards work between the IETF and the International Telecommunication Union (ITU), the result was that the directory problem fell between the cracks of the two organizations.

This dilemma was finally addressed with the creation of the Voice over IP Forum (VoIP). VoIP is now a working group of the International Multimedia Telecommunications Consortium (IMTC), which is an industry group of the ITU. Because VoIP is in the IMTC, it does not have to follow ITU rules and protocols. In the past several months directories and CMA have been discussed at several VoIP meeting, with a large percentage of discussion on CMA.

The discussion around CMA and directories has had considerable overlap. A directory-centric approach provides a simple solution to the interoperability problem among the directories of existing applications, without calling for a change in the basic design of existing directories. CMA calls for an entirely different approach be taken to user location than existing directory designs. Because of the complexity of CMA and the patent that VocalTec currently holds on the design, it is unlikely that it will be ready to be used in the near future. Because of this, one attractive solution would be to implement a directory-centric solution now to address existing interoperability problems and to later migrate that solution to a CMA design.

### 8.3 ITU and IETF Design Philosophies

This section provides a brief discussion of the different design philosophies of the IETF and ITU, which provides a basis for the economic discussion in Section 8.4. On a technical level, one difference between the International Telecommunication Union (ITU) standards and the (Internet Engineering Task Force) IETF standards is the assumption of which communications "stack" they are running on. A communications stack specifies both the interface and the services to a network, or a range of networks. For communications to be successful then both sides must use the services from the same stack or they must use a gateway, which translates between protocols. The stack from the ITU is called the OSI stack, for Open Systems Interconnection. The "stack" that comes from the IETF is most often called the TCP/IP protocol suite (Transfer Control Protocol/Internet Protocol). It is not entirely accurate to think of each stack being an integrated whole. Not all stacks will contain all services, only those services that are used are needed at both sides.

The different terms "stack" and "suite" are used for a reason. That reason is to reflect the different design philosophies of the IETF and the ITU. The word stack reflects a rigid hierarchy, in which each standard is dependent on all those below it. The word protocol "suite" reflects a more loose relationship between protocols. This allows the designer to "pick and choose" which protocols they want in their suite, and only get those that they need. This reflects the IETF's design philosophy based on flexibility.

This difference in design philosophies has had a very significant effect on the standards proposed by the ITU and the IETF. Although attempts have been made to try to integrate the ITU and IETF standards, it is obvious that they have differences that are based on the design philosophies of each of the two groups. The ITU standards are monolithic, and do not specify interfaces between important layers. Only the end-to-end interfaces are specified. This creates a situation in which the standards must be taken in an "all or nothing" basis, making it difficult for someone implementing the standards to pick and choose what they want to implement. To be compliant with ITU standards and to receive "certification" they must comply with all mandatory parts, which is the majority of the standard.

In my own opinion, the ITU standards are far more complex and difficult to understand than they need to be for the Internet. This complexity comes from the fact that the standards were first designed for the PSTN, and then "patched" to the Internet. A design for the Internet could have left off a large portion of the complexity. The H.323 series does better at avoiding these problems than most ITU standards, and T.120 standards does an even better job than H.323, but they still do not approach the simplicity and modularity of IETF standards.

The IETF conferencing standards are based on a model called Application Level Framing (ALF) [Cl90], which leaves many decisions up to the application, independent of the underlying physical network. This creates standards that are composed of many smaller components that are much easier to understand, and their modularity contributes to their flexibility. The one criticism of the IETF standards is that they do not give the appearance of providing a system that has been fully completed. This is because they are designed primarily to focus on the use of multicast for conferencing, which is still in the process of development. The conferencing directory proposed in this paper could be used by both the IETF or the ITU because it is based on the X.500 directory. This is because the X.500 directory is available both through both the OSI stack (through Directory Access Protocol or DAP) and the TCP/IP suite (through LDAP).

This design philosophy of a rigid hierarchy has made sense historically for the ITU for quite some time because that was the best way for the PSTN to be built. Put in economic terms this was because there were economies of scale in the underlying technology of the PSTN. In computer science terms this translates into a fairly "closed" system because the economies of scale from the technology at the time did not make it very useful for outsiders to contribute to the network. The way a system is made closed is either by making interfaces very complicated or by not specifying interfaces. The term "closed" means that you cannot see the interfaces. This makes each component of a system dependent on the one "below" it, so that you must choose to take the whole system or none at all.

In all cases, users at both ends must have matching protocols (or gateways). But a flexible "suite" will allow users to select only the protocols *for the services that they use* and *will allow them to choose protocols from different vendors*. Not only does this promote competition, but it allows for an enormous diversity of applications to be generated. In a hierarchical stack, only the provider of the stack can add protocols and services for new applications, which gives them control over the application. The design philosophy of the IETF has been one which allows for flexibility by openly specifying interfaces and separating components into the smallest units necessary.

The reason for this design philosophy was that the technology of the IETF was small components of software. Put in economic terms, if a single component of software is small enough, then there are absolutely no economies of scale in writing a small component of software. In fact there are significant diseconomies of scale that result in the design of a very small software component. As an extreme example, one person is better than two in writing one line of code. This economic principle may not hold in combining and reselling components (as Microsoft demonstrates), but it does hold in the design of small individual components. The reason for the diseconomies of scale in making small software components is that it does not matter who makes a component if it functions according to its specification. Anyone with knowledge and a compiler can write the code for the software, so there the optimal economy of scale of a "small" component is one person. It is based on this fact that object oriented program has emerged, and all software development is modularized to allow one person to focus on modules or components of code. Because the Internet was such a large system composed of software, the only way to design it was through a flexible philosophy that separated components into the smallest effective unit and specified their interfaces. These different design philosophies are reflected in the different conferencing standards of the ITU and the IETF.

### 8.4 Interoperability Between the Internet and the PSTN: Economics of Decoupled Services and Network

The most significant effect that multimedia conferencing on the Internet has on existing market structure is that it "decouples" services from the network [NI96]. Because conferencing directory services play a large role in that decoupling, it is important to consider the economic effects that this might have. The term "decoupling" means that the network provider no longer has control of the majority of the services you receive. The effect of complete decoupling is that third parties can sell services rather than the service provider [Go96]. In the case of Internet conferencing, many of these third party services come from software, which either is free or has a very low price compared to the carrier's price. The carrier is only able to receive revenue on the very limited services that it does provide, and is not able to add revenue for each additional service. For example, call waiting can be included as a feature in conferencing software at a price of zero, and since it can be added without adding any new requirements to the carrier, the carrier cannot charge for this new service.

## *8.5 Internet Telephony and Regulation*

It is likely that one of the most significant effects of Internet telephony will be on regulatory policy. Internet telephony represents the convergence of the telecommunications industry, which is heavily regulated, and the computing industry, which is almost entirely unregulated. Because of this, Internet telephony is bringing major regulatory issues to the surface. In 1996, a group of long distance resellers called America's Carriers Telecommunications Association (ACTA) filed a petition before the FCC requesting a ban on Internet telephony [Ac96]. The FCC has announced that it will not act on the petition at this time.

Several Local Exchange Carriers used the discussion of Internet telephony to bring up the issue of local access fees. Local access fees are the per-minute fees that the interexchange carriers, such as AT&T pay for each call (about 2.5 cents/min on each end). The fees cover the cost of the local access line, but also serve to provide a subsidy which goes into the universal service fund. Because Internet Service Providers (ISP's) are classified as "Enhanced Service Providers," they are exempt from paying the local access fee. The local carriers argument was that since the Internet can now carry voice, ISP's should be required to pay local access fees, which would raise Internet access rates by $1.50/hour. The FCC has since put out a request for comments on the issue of access fees for ISP's.

The recent debates over the ACTA petition and access fees for ISP's represent only a few of the many regulatory issues that the Internet is bringing to the forefront. Regulatory bodies in countries from across the world are also facing similar regulatory issues. One of the most significant issues is that of international tariffs. Many countries in the world have high tariffs on all international calls, but Internet traffic is not charged the high tariffs. This has created significant opportunities for arbitrage by transmitting voice over the Internet.

The overall effect the Internet is having on regulation of telecommunications seems to be to create pressure to speed up the existing trend toward deregulation. Another view of this trend might be that Internet telephony might cause an increase in the regulation of the Internet, which has been virtually unregulated. This is what one might expect that when a regulated system converges with an unregulated system: they are likely to reach a level of regulation between the two systems.

One key question for policy makers to consider is whether one system or the other might be dominant, or whether a new common system emerge which has different features than either system. In considering the basic architecture of the PSTN vs. the Internet, one can make an argument that the level of regulation of each system suited it. If the Internet were to dominate a global network, then a strong argument could be made that the network should be largely unregulated based on the technology. This is because the basic technology of the Internet is largely decentralized and is not likely to have the same economies of scale as the PSTN. This conclusion is an assumption and would need to be studied further to verify its validity.

# 9. Conclusion and Suggestions for Future Work

This paper presented a design and implementation of group and user location for conferencing in LDAP. This design shows that LDAP provides all the functionality needed for conferencing directories. Existing conferencing application providers can utilize the design along with their existing directory system to maintain backwards compatibility. The design achieves scalability by relying on fine-grained partitioning rather than replication and caching to distribute load. There is still a considerable amount of work to be done to provide conferencing directories compatible with existing standards.

There are a number of specifications that might be needed to gain the full benefit of conferencing directories. A partial list of specifications for record types and organization is as follows:

**User Location:** Need formal specification of attributes related to user location to include in user records.

**Group Information**: Need formal specifications for attributes for a group information record. Need specifications to translate T.124's Conference Query command, SAP/SDP and IRC Channels to subclasses of group information records.

**Group Membership**: Need formal specifications for a common set of attributes for group membership. Need specifications to translate the Conference Roster command in T.124, SDES packets in RTCP, and IRC members to subclasses of group membership records

**Organizational Schema**: Need formal specification of where in the X.500 tree organizations should store references to their static groups and aliases. Need location for aliases and static group information can be found. Need an independent organization (Internic, ISI, etc) to store primary copy of redirection references to aliases and static group information.

## 10. References


[Ac96] America's Carrier's Telecommunications Association. "ACTA Petition Before the FCC." Mar. 1996. http://www.fcc.gov/Bureaus/Common_Carrier/Other/actapet.html

[Bo96] Bormann, C., Ott, J., Reichert C. "Simple Conference Control Protocol." Internet Draft. June, 1996. draft-ietf-mmusic-sccp-00.txt.

[Br97] Braden, R. et al. Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification." June 14, 1997. Internet Draft. ftp://ftp.ietf.org/internet-drafts/draft-ietf-rsvp-spec-16.txt

[Cl90] Clark, David D. and David L. Tennenhouse. "Architectural Considerations for a New Generation of Protocols." SIGCOMM '90, pages 200-208, Philadelphia, Sep 1990. ACM.

[Fe97] Fenner, W. "Internet Group Management Protocol, Version 2." Internet Draft. January 20, 1997. http://ds.internic.net/internet-drafts/draft-ietf-idmr-igmp-v2-06.txt

[Fl95] Floyd, Sally, Van Jacobson, Steve McCanne, Ching-Gung Liu and Lixia Zhang. "A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing." ACM SIGCOMM 95, Cambridge Mass, Sep 1995.

[Fi96] Fielding, R. et al. "Hypertext Transfer Protocol 1.1." Internet Draft. August 12, 1996. ftp://ftp.ietf.org/internet-drafts/draft-ietf-http-v11-spec-07.txt

[Ge96] Genovese T. "A Common Schema for the Internet White Pages Service." Internet Draft. November, 11, 1996. ftp://ftp.ietf.org/internet-drafts/draft-ietf-ids-iwps-schema-spec-02.txt

[Go96] Gong, Jiong and Srinagesh, Padmanabhan. "An Economic Analysis of Network Architectures" IEEE Network. March/April 1996. (a paper presenting a similar argument can be found at http://www.press.umich.edu:80/jep/works/GongEconLa.html)

[H.225] ITU Recommendation H.225 "Line Transmission of Non-Telephone Signals." 1996.

[H.245] ITU Recommendation H.245 "Multimedia System Control." 1996.

[H.323] ITU Recommendation H.323 "Visual Telephone Systems and Equipment for LANs without guaranteed QoS." 1996.

[Ha96a] Handley, M, J. Crowcroft, C. Bormann. "The Internet Multimedia Conferencing Architecture." Feb. 22, 1996. Work in Progress.

[Ha96b] Handley M., V. Jacobson "SDP: Session Description Protocol." Internet Draft. draft-ietf-mmusic-sdp-02.ps, Febuary 27, 1996.

[Ha96c]Handley, M. "SAP: Session Announcement Protocol." Internet Draft. June 3, 1996. Work in progress. draft-ietf-mmusic-sap-00.ps.

[Ha96d] Handley, M. Schooler, E. "Session Invitation Protocol." Internet Draft. Feb, 1996. draft-ietf-mmusic-sip-00.ps

[ITU-TS] ITU-TS Group. "Draft ODP Trading Function." July, 1994.

[Jo94] Johannsen, T, G. Mansfield, M. Kosters, S. Sataluri, "Representing IP Information in the X.500 Directory", 03/25/1994. http://andrew2.andrew.cmu.edu/rfc/rfc1608.html

[Ma93] Malamud C., M. Rose. "Principles of Operation for the TPC.INT Subdomain: General Principles and Policy." RFC 1530. October, 1993.

[Mo86] Mockapetris, P. "Domain Names -Concepts and Facilities." RFC1034. Nov, 1987. http://andrew2.andrew.cmu.edu/rfc/rfc1034.html

[Mo86] Mockapetris, P. "Domain System Changes and Observations." RFC973. January, 1986.

[Ne93] Neuman, B Clifford, Steven Argart. "Prospero: A Base for Building Infrastructure." Proc. INET '93. August, 1993.

[NI96] NII 2000 Steering Committee, National Research Council. "The Unpredictable Certainty: Information Infrastructure Through 2000." National Academy Press Washington, D.C. 1996.

[Oi93] Oikarinen J., D. Reed, "Internet Relay Chat Protocol." RFC1459, 05/26/1993. http://andrew2.andrew.cmu.edu/rfc/rfc1459.html

[Pe97] Petrack , Scott. "Call Management Agent System Requirements Function Architecture and Protocol." VoIP Contribution 97-10. Jan. 1997.

[Scho95] Schooler, E. "A Multicast User Directory Service for Synchronous Rendezvous" Master's Thesis. May 5, 1995.

[Schu96a] Schultzrinne H., Casner, Fredrickson, V. Jacobson. "RTP: A Transport Protocol for Real-Time Applications" RFC 1889. February 1, 1996.

[Schu96b] Schultzrinne, H., "Personal Mobility for Multimedia Services in the Internet" IMDS'96, March 4-6, 1996. Ftp://ftp.fokus.gmd.de/pub/step/papers/Schu9603:Personal.ps.gz.

[Schu96c] Schultzrinne, H., "Simple Conference Invitation Protocol" Internet Draft draft-ietf-mmusic-scip-00.txt. February 22, 1996.

[Se95] Sears, Andrew L. "The Effect of Internet Telephony on the Long Distance Voice Market." Paper, presentation." European Competition, Law and Regulation, London. Dec, 1996

[Se96a] Sears, Andrew L. and Kevin Savetz, "FAQs How Do I Use the Internet as a Telephone?" 1996.

[Se96b] Sears, Andrew, Phil Mutooni, Rob Cheng, and others. "A Cost Model of Internet Telephony for Regulatory Decision Making." FCC Submission. May, 1996.

[Se97a] Sears, Andrew. "A Scaleable Directory Schema in LDAP for Integrated Conferencing Services." INET 97. July, 1997. Kuala Lumpur, Malaysia.

[Se97b] Sears, Andrew. "The Economics of Internet Access in the United States: Strategic and Regulatory Implications." INET 97. July, 1997. Kuala Lumpur, Malaysia.

[T.120] ITU Recommendation T.120 – Data Protocols for Multimedia Conferencing. 1995

[T.124] ITU Recommendation T.124 – Generic Conference Control. 1995.

[Vi96] Vixie, Paul, Susan Thomson, Yakov Rekhter, Jim Bound. "Dynamic Updates in the Domain Name System (DNS UPDATE)" March 1996. Internet Draft. Work in Progress. draft-ietf-dnsind-dynDNS-09.txt

[Wa96] Wahl M., T. Howes, S. Kille. "Lightweight Directory Access Protocol (v3)" Proposed Internet Standard. June 3, 1996. ftp://ftp.ietf.org/internet-drafts/draft-ietf-asid-ldapv3-protocol-01.txt

[We92] Weider C. , J. Reynolds, S. Heker. "Technical Overview of Directory Services: Using the X.500 Protocol" RFC 1309. March 1992. http://andrew2.andrew.cmu.edu/rfc/rfc1309.html

[Wi96] Williams, K "User Location Service" Internet Draft. February 1996. http://rpcp.mit.edu/~itel/uls.html

[X.500] ITU Recommendation X.500. "The Directory: Overview of Concepts, Models, and Services.

[Ya96] Yaacovi Y., K. Settle. "Lightweight Directory Access Protocol: Extensions for Dynamic Directory Services" Internet Draft Standard. June 13, 1996. ftp://ftp.ietf.org/internet-drafts/draft-ietf-asid-ldapv3ext-00.txt