



Integrating Internet Telephony Services

The Cinema system uses a SIP-based architecture to integrate Internet telephony services with an existing PBX infrastructure.

**Wenyu Jiang,
Jonathan Lennox,
Sankaran Narayanan,
Henning Schulzrinne,
Kundan Singh,
and Xiaotao Wu**
Columbia University

Cost savings and the ease of developing and adding new services have motivated great interest in Internet telephony, which integrates services provided by the Internet with the public switched telephone network (PSTN). Internet telephony relies on several protocols, including the real-time transport protocol (RTP)¹ for multimedia data transport and the session initiation protocol (SIP)² or H.323³ for establishing and controlling sessions. SIP can integrate with other Internet services, such as e-mail, the Web, voice mail, instant messaging, conference calling, and multimedia collaboration.

We have implemented a SIP-based software suite called the Columbia Internet extensible multimedia architecture (Cinema), which we installed and integrated with the existing private branch exchange (PBX) infrastructure in the computer science department at Columbia University. The Cinema environment

provides interoperability with the PSTN, programmable Internet telephony services, and IP-based voice mail. It also integrates Web access and e-mail for unified messaging and supports multiparty multimedia conferencing. The setup lets us extend our PBX capacity and will eventually let us replace it while keeping our existing phone numbers. It also provides an environment in which we can easily add new services and features, including interoperation with existing multimedia tools, e-mail access from standard telephones, network appliance control, and instant messaging support. We believe Cinema can be readily adopted by other organizations.

Architecture

For a standard Internet audio call, a participant must know the audio codecs the other participant supports and the IP address and port number where the other participant wants to receive the audio

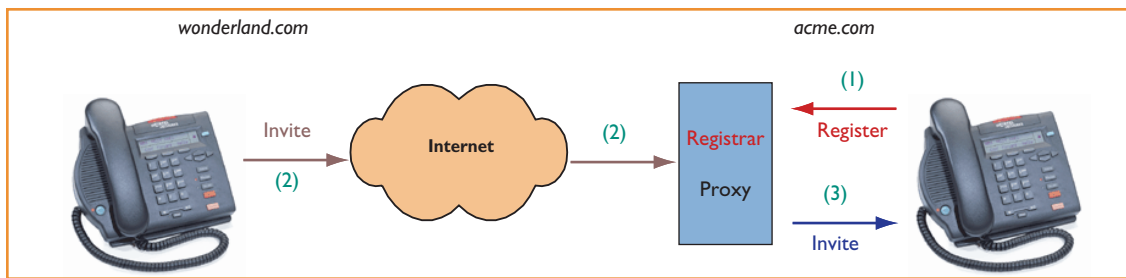


Figure 1. SIP call flow. (1) When Bob's phone powers up, it registers with the server `acme.com`, indicating that Bob can be reached at `p42.acme.com`. (2) To initiate the Internet call to `sip:bob@acme.com`, Alice's SIP phone sends an invite message to the `acme.com` server, which knows the locations at which Bob can be reached. (3) In redirect mode, the server will return Bob's location to Alice, or in proxy mode, the server will forward the invite message itself to try to contact Bob at his current location.

packets. The problem with this model is that IP addresses are hard to remember and can easily change if the user changes locations or machines. SIP facilitates user mobility by using high-level addresses of the form `user@domain`. For instance, a user can call `bob@acme.com` regardless of what communication device, IP address, or phone number Bob uses. SIP registrar servers maintain user locations, and the user's communication devices register with the registrar servers periodically by providing their current addresses.

Figure 1 shows the steps involved when user Bob wants to call another user, Alice. A forking proxy can forward the call request to multiple locations sequentially or in parallel. The first phone that is picked up gets the call, and the rest of the phones stop ringing. SIP calls can also use "tel" URLs that identify E.164 telephone numbers such as `tel:+12125551234`. The session description protocol (SDP), carried in SIP requests and responses, describes the list of supported audio and video codecs and the transport addresses to receive them.⁴

Cinema Components

Figure 2 (next page) shows the Cinema architecture and the interactions among the system components, which consist of a set of SIP-based servers that provide a pathway to a post-PBX era of communications. The SIP server (`sipd`) and the SQL database form the core of the Cinema infrastructure. `Sipd` is a proxy, redirect, and registration server that helps route calls to a user's correct location based on information stored in the SQL database. The system assigns each user a *canonical user identifier* of the form, `user@domain`. Generally, Cinema assigns users their e-mail addresses as SIP identifiers, but it can also operate in portal mode, in which the system creates a new identity specifically for SIP calls.

Cinema uses a Web-based interface for managing user accounts, voice mail, and conference calls. Users can alter their passwords, authentication mechanisms, or voice mail notification format, and they can use the interface to update their network addresses and the phone numbers at which they can be reached.

As shown in Figure 2, the system includes several specialized servers. Cinema's general-purpose streaming media server, `rtspd`, delivers announcements, conference recordings, and voice mail messages stored there by the system's unified messaging server, `sipum`, which is essentially a centralized voice mail system.⁵ (Users can retrieve voice messages using a real-time signaling protocol [RTSP] client, such as Apple's QuickTime.) Finally, `sipconf` is Cinema's centralized audio and videoconferencing server.⁶

The `sip323` translator is a signaling gateway between SIP and H.323 – the International Telecommunication Union's (ITU-T) standard for multimedia conferencing over any packet-based network. In Cinema, `sip323` integrates popular H.323 clients, such as Microsoft NetMeeting, into a SIP infrastructure. A Cisco 2600 router with SIP and PSTN capability acts as a gateway between the departmental telephone switch (PBX) and the department's local area network (LAN).

Call Handling

Users make calls with IP-based SIP user agents (UAs), which can be based on either hardware (such as our Ethernet phone) or software (such as our `sipc` software system that runs on PCs). Figure 3 (page 67) shows the steps involved in processing a call when, for example, `sip:alice@wonderland.com` calls `sip:Bob.Wilson@cs.columbia.edu`.

Through DNS SRV⁷ service location records, Alice finds out that the server `erlang.cs.columbia.edu` serves the `cs.columbia.edu` domain,

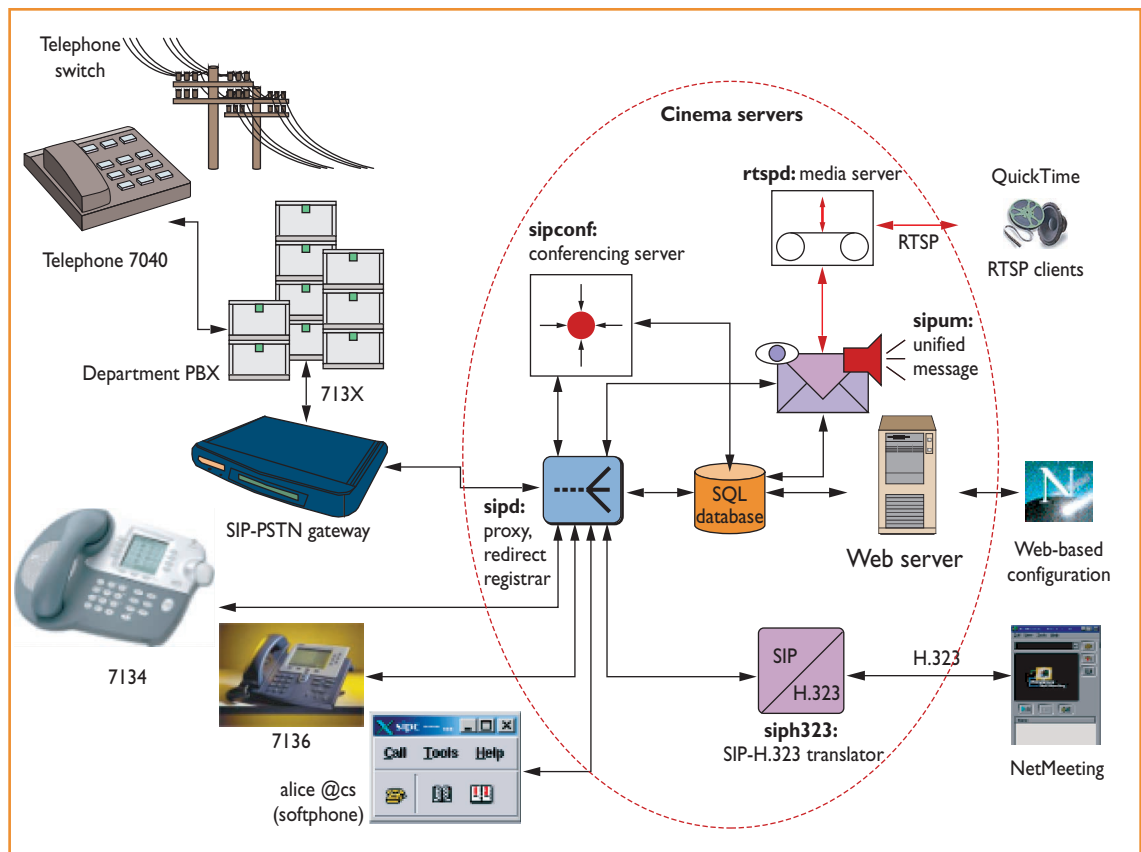


Figure 2. Cinema architecture. The SIP server (sipd) is a proxy, redirect, and registration server that works with a SQL database to route calls to a user's correct location. The system also includes servers and gateways to various services.

and sends the request to it. On receiving the message, the server transforms the callee address to a canonical user identifier for database lookup by a process called *canonicalization*. In this case, Cinema canonicalizes erlang.cs.columbia.edu to cs.columbia.edu by matching the domain portion of the request URI against a regular expression. For example, the line `(.*cs.columbia.edu) | (128.59.1[4-9].*)` maps all host names and IP addresses in the range 128.59.14.0 – 128.59.19.255 to the canonical server address cs.columbia.edu. If the canonicalized host name does not match, the server operates as an out-bound proxy and routes the unprocessed request to the SIP server for the domain.

In Figure 3, we assume that the system can reach user Bob with the canonical identifier sip:bob@cs.columbia.edu in several ways, including bob, Bob.Wilson, Bob.V.Wilson, or through a phone number such as 7134. To map such non-standard representations to a canonical identifier, the system performs the four types of mappings shown in Figure 3: SQL alias, system alias, name mapping, and dial plan.

First, the system checks the database for the presence of a user's alias and, if one is present, resolves it to its canonical identifier. Otherwise, the server checks the system's e-mail aliases file, which typically records functional aliases like "postmaster" or "Webmaster." If no match is found, the name-mapper function tries to deduce a username by comparing various combinations of the first and last name recorded in the password entry against the system password file. Finally, if the user identifier looks like a telephone number, such as sip:7018@cs.columbia.edu or tel:7018, the system performs dial plan transformations. If none of the above rules locates a match, the system uses the user identifier unchanged.

In Figure 3, canonicalization results in the user name bob@cs.columbia.edu, and the SIP server retrieves the corresponding user, contact, and policy information from the database. If Bob requires caller authentication, the SIP server tries to authenticate Alice. The system then executes Bob's policy, which might, for example, demand that a particular calling user be refused or redirected. In addition, Bob's policy might involve executing

programmable scripts that dynamically decide where Bob wants to be reached. If the server determines that Bob's current policy allows Alice's call, it contacts all of Bob's registered locations. Bob's SIP phone rings, he picks up the handset, and starts talking to Alice. Either participant can terminate the call when they finish talking.

Unlike PSTN switches, SIP servers do not normally store call state. They are responsible only for forwarding signaling messages, such as call initiation and termination requests. While the call initiation request goes through the SIP server, call termination requests might be directly exchanged by the two user agents without any intervening SIP server. However, the server can insist on being in the call path for all further messages using the SIP record-route option, which is useful for call logging and accounting.²

PSTN Interoperation

Because telephone numbers identify PSTN subscribers, PSTN callers can reach SIP users by calling the gateway with a telephone number in the range assigned to the gateway's trunk line. The telephone number range is assigned either by the local PBX, if there is one, or by the local telephone company. For example, our PBX assigns extensions in the range 7130-7139 to the Cisco 2600 gateway, and anybody who dials (212) 939-7134 reaches the gateway at extension 7134. On the gateway, the system defines a VoIP call-leg specifier (called a *dial peer*) that determines a mapping to the SIP server's IP address, the type of codec that must be used in that call leg, and other parameters. In the following example, `destination-pattern` defines the range of numbers assigned to SIP phones, and `session target` gives the SIP server's IP address.

```
dial-peer voice 1 voip
  preference 1
  destination-pattern 713[0-9]
  voice class codec 1
  session protocol sipv2
  session target ipv4:128.59.19.62
```

When a call comes in from the PSTN, the gateway can react in one of two modes: direct inward dialing (DID) or no-DID. In DID mode, the incoming trunk delivers the destination extension to the PBX or gateway. The SIP server maintains a map, called a *dial plan*, between the telephone number and the user identifier. In no-DID mode, the gateway prompts the caller with a second dial tone. After

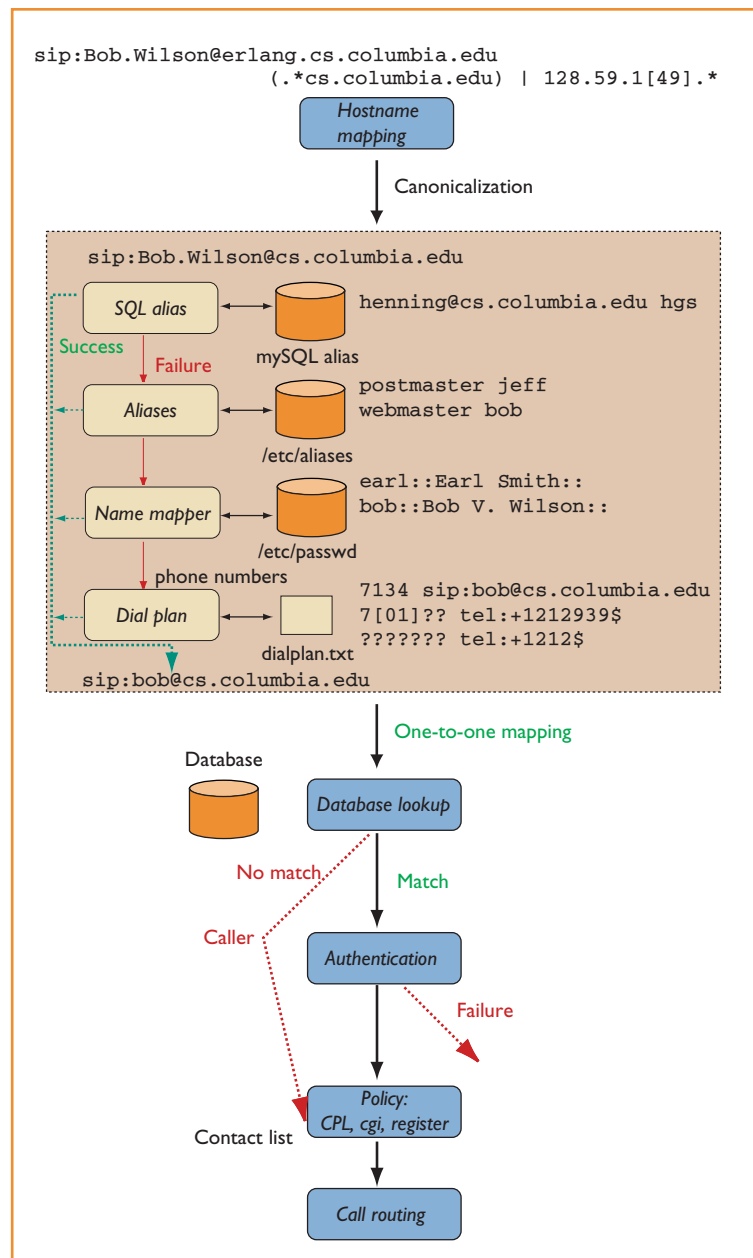


Figure 3. Call-handling procedures in the Cinema system. (1) Callee look-up (Bob) employs hostname canonicalization, user id canonicalization, and user database look-up; (2) Caller authentication allows only authorized persons to reach the callee; (3) Call routing forwards the call to the callee's current location or, if a policy is stated, to where the callee wishes to be reached.

the caller dials a new extension, the gateway captures the call and forwards it to the SIP server. DID mode allows simpler dialing from PSTN, but no-DID mode can support more users. In the reverse direction, when a SIP user dials a telephone number, the SIP server transforms a telephone number such as `sip:9397040@cs.columbia.edu` to a telephone subscriber such as `tel:+12129397040`.

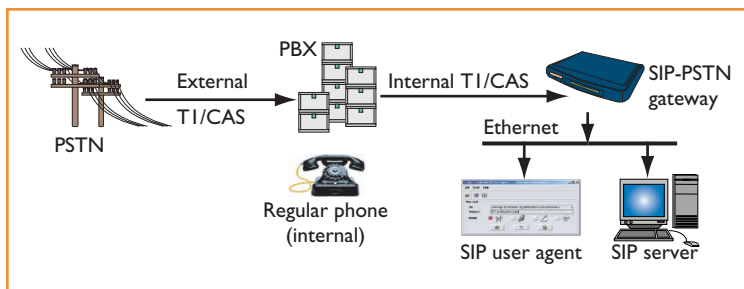


Figure 4. A standard PBX configuration. The arrows illustrate how a call goes from PSTN to a SIP user agent (UA) via the PBX, the gateway, and the SIP server.

Translating SIP-to-PSTN Calls

Most PBX systems allow the same number to be dialed in multiple ways. For example, callers can use a four-digit extension, a local phone number, or a global number that includes the country code. The SIP server must therefore be able to unify numerous variations into a single number that can be used to select the appropriate gateway.

The following sample dial plan illustrates how this model operates. It maps both 4-digit extensions and 7-digit local numbers to a canonical format that includes the country code, area code, and local number. The SIP server substitutes the matched string in the left column for the \$ symbol; ? matches a single digit, and * matches any digit string. Any numbers in the round brackets, such as the (8) and (011) dialing prefixes in the example below, are removed during translation.

```
# Intradepartment calls
7[01]??      tel:+1212939$
# Local (same area code) calls
???????     tel:+1212$
# '8' is used for long distance
(8)???????? tel:+1$
# International numbers
(011)*      tel:+$
```

The SIP server then locates the appropriate gateway to route the call to the PSTN. For organizations using few gateways, the SIP server uses a static mapping table, called a *gateway map*, to select a gateway for each user or group. The gateway map also specifies which user groups can make certain types of gateway calls, such as international calls. If networks of IP telephony gateways are deployed, more complex routing protocols, such as telephony routing over IP (TRIP),⁸ might become essential. TRIP lets the system route calls to the optimal gateway – usually the one closest to the destination.

Connecting to the PBX

For centralized phone management, many corporations and universities use PBXs to consolidate external trunk lines and provide voice mail services. Our department has a Nortel Meridian Option 11C PBX with two T1 trunk lines.⁹ The external T1 line connects to the local central office, and the internal T1 line connects to the PSTN-IP gateway, as shown in Figure 4. With this setup, a user can make IP phone calls from either an analog phone (whether inside or outside the department) or a SIP user agent. The integration of the PBX and PSTN-IP gateway is usually a trial-and-error process, largely due to the proprietary and legacy nature of PBX systems.

Line type. The T1 line type can be either channelized or primary rate ISDN (PRI).⁹ T1s differ mainly in the signaling protocol they use. Channelized T1s support 24 64-Kbps voice channels and use channel-associated signaling (CAS) – a form of in-band signaling in which some bits in each voice channel are robbed for signaling. In contrast, PRI reserves one channel for an out-of-band signaling method known as common-channel signaling (CCS). PRI supports 23 voice channels plus one signaling channel. Depending on the PBX model, channelized T1 might not provide advanced features such as caller ID, and hence PRI service might require additional signaling hardware in the PBX's T1 line-card. We use channelized T1 in our PBX for both T1 lines.

Line characteristics. The line coding determines how the bits are coded on the wire. We recommend binary 8-zero substitution (B8ZS) because it provides 64 Kbps for each voice channel, whereas alternate mark inversion (AMI) steals one out of every eight bits (leading to a 56-Kbps channel), which degrades voice quality. Line coding is not always independent of line type. AMI cannot be chosen with PRI, for instance, because PRI requires 64-Kbps channels.

Trunk type. The most common choices for trunk type include DID and tie line. The name “tie” comes from the fact that this type of trunk is a bidirectional trunk line that ties two nodes together. We recommend using a tie line because it allows both DID (incoming) and outgoing calls. The channel type is also important. It can be data, voice-only, or data and voice, but it must be set to voice-only for channelized T1. Otherwise, IP-to-PSTN calls might fail as the PBX treats a call as data transmission.

Related Work in Telephony Integration

Several companies, such as Net2Phone and MediaRing, provide PC-to-PC and PC-to-phone calls. Their objective is mainly to provide low-cost call service to the PSTN from the public Internet. We initially intended Cinema to minimize telephone infrastructure and service costs for an organization, but our architecture is well suited as an Internet telephony infrastructure within an organization. We can configure Cinema to carry calls between campuses or branch offices over IP with virtually no added cost.

Several multimedia conferencing products use SIP or H.323 for signaling. These include MeetingPoint from CUseeMe Networks (www.cuseeme.com), Sametime from Lotus (www.lotus.com/home.nsf/welcome/sametime), and GnomeMeeting (www.gnomemeeting.org) from the Linux community. Our system can provide services beyond standard videoconferencing and can actually incorporate these tools as long as they are standards-compliant.

There is a fair amount of early voice-

messaging work — in particular, Xerox PARC's Etherphone — but none of it addresses the integration of Internet telephony with voice-messaging systems. Several vendors offer SIP proxy servers and user agents that we can use in the Cinema infrastructure as well. The SIP standard simplifies integration among different vendors' products. The Cinema architecture facilitates such integration with readily available features; new components can be added as needed.

Access permissions. Nortel Meridian PBXs use multiple priority levels called network class of service (NCOS) for controlling user privileges. A caller's NCOS determines the type of call users are allowed to make. For instance, a low NCOS value indicates a low privilege level, and might allow only internal phone calls. For incoming calls, the system checks the caller's NCOS, and the call goes through only if the caller's NCOS permission is no less than the callee's. (The callee is not the internal T1 line but a virtual entity.) To enable incoming calls to SIP users, we configure the callee's NCOS value to match the caller.

Security Issues

Because user registrations must be authenticated to prevent unauthorized users from directing calls to themselves or elsewhere, our system uses digest authentication. This means that the system will always verify a shared secret between the server and the client via challenge-response prior to allowing access. Currently, our PSTN-IP gateway does not have user authentication and authorization capabilities, so we delegate this functionality to the SIP proxy server to ensure that only authorized users can make calls.

Our gateway is a Cisco 2600 router, which runs the internet network operating system (IOS) software. We make use of the IOS access control lists (ACL) to accept SIP requests only from the proxy server. This prevents users from bypassing the proxy and using the gateway directly, although the gateway is configured to allow media packets from any host. The system forwards calls from the gateway to the SIP server. This configuration, along with the SIP record-route mechanism,² forces subsequent requests within a call to traverse the same signaling path

and allows call logging and billing services to be part of the SIP server.

Additional Services

An ideal telephony system would provide support for intelligent call-handling scenarios. For example, users might want other users to be able to reach them at their office phone during office hours and at their home phones after hours, but they might not want telemarketers to get through at either location. Users can achieve this level of customizability with the Cinema system by uploading code that alters the SIP server's behavior based on call priority or caller identification. Several different mechanisms can achieve this with SIP, including the XML-based call-processing language (CPL)^{10,11} or the SIP common gateway interface,¹² which behaves much like standard HTTP CGI.

Our SIP user agent, *sipc*, simplifies creation and deployment of SIP CGI and CPL scripts by providing scripting engines for these languages. In addition to scripting, however, user agents need services like call signaling and media transmission control between endpoints. Because SIP CGI and CPL are not designed for such endpoint services, we are currently developing a new service script language specifically for supporting endpoint services like direct media control and direct user interaction.

Voice mail is as important for Internet telephony as it is for the PSTN. Unlike answering machines, voice mail systems enable centralized user account management and configuration, and they make it easy for users who have multiple phones. Internet-based voice mail integrates easily with other IP services. We use SIP for signaling and RTSP for storage and retrieval of voice messages.⁵ The system can list user messages on a Web

example.com

```
_sip._udp      0 40 a.example.com
                0 40 b.example.com
                0 20 c.example.com
                1  0 d.backup .com
```

Figure 5. DNS SRV resource record. This record orders servers by priority and weight for each service and domain.

page where they can be played by a mouse-click using any of several media players.

Conference Calling

Another important telephony service is multi-party conferencing, which is provided in the PSTN by conference bridges. Our Internet telephony environment employs a SIP conference server⁶ that uses audio and video capabilities for unicast and multicast conferences. Every conference operates using a SIP address, such as `staffmeet@cs.columbia.edu`, which users dial to join the conference. The dialplan facility in the SIP server helps map telephone numbers to conference addresses, which facilitates participation by PSTN users.

The Web interface offers access to attributes such as conference identifier, duration, schedule, authentication mechanism, conference size, and media type. The system can restrict the participant list by defining specific capabilities for different sets of participants. Participants might want a conference in which anybody could listen, for instance, but only users within the domain `cs.columbia.edu` could send media.

Interactive Voice Response

An Internet telephony system should also provide services such as authentication or listening to and deleting voice mail for users who are limited to supplying input using the telephone keypad and spoken audio. The VoiceXML markup language can help in creating audio dialogs that feature synthesized speech, digitized audio, input speech, dual-tone multifrequency (DTMF) recognition, and audio recording for telephony applications. The separation of user interaction code (such as DTMF and speech recognition) from the service logic (phone-based e-mail, weather reports, and so on) facilitates the development of several interactive voice-response services.

We built a SIP-based VoiceXML browser that lets telephone users interact with Cinema components. This component operates between the SIP-

PSTN gateway and the voice mail or conferencing service in our system. VoiceXML defines a dialog specification for services such as voice mail access. Our browser fetches this dialog from a Web server using HTTP and interprets it in a form suitable for PSTN users.

Network Appliance Control

Beyond providing multimedia communication services, we also developed a SIP/X10 gateway that provides network appliance control. X10 is a protocol that involves sending specially coded low-voltage signals that are superimposed over the electrical wiring. Our gateway accepts SIP requests and translates the requests into X10 signals, which are then transmitted to the controlled devices to change the device status to on or off. Each device attached to the gateway has a SIP address, such as `sip:lamp@cs.columbia.edu`. Users who subscribe to the SIP address can be notified by the device status change.

Instant Messaging and Presence

By combining presence and Internet telephony, providers can improve services – reducing the number of failed call attempts or automatic redirects to voice mail, for example. Users who subscribe to each other can be notified of changes in state (online, offline, busy, and so on) and can send each other short instant messages. SIP-based messaging and presence provides a standard way to send instant messages and form “buddy lists.” We can also couple presence to phone status, marking the user as busy when in a phone call.

Our SIP server, `sipd`, has a built-in centralized presence server that stores and conveys presence status. In this model, to convey her interest in Bob’s presence status, Alice subscribes to a server that tracks subscriptions for each user. When Bob comes online, his user agent registers with the SIP server and provides his contact address. The server then uses this registration information to notify Alice about Bob’s status asynchronously. The system displays buddy status as part of the user’s address book, and users can exchange instant messages with their buddies, directly or through a proxy server, using the `SIP message` method.

Advanced Services

In addition to the services offered by Cinema servers, intelligent SIP user agents can provide advanced services beyond basic call establishment. The `sipc` user agent supports instant mes-

saging and can retrieve presence information about buddies from a SIP presence server such as sipd. The agent also facilitates application sharing using the virtual network computing (VNC) tool,¹⁴ as well as shared Web browsing, which lets a group of people visit the same set of Web sites in a session. Here, user agents use the SIP message feature to exchange URLs¹³ with each other on the basis of which sites a particular user is currently visiting. These services help enhance our communications infrastructure to create a collaborative work environment.

Scalability

The distributed nature of our system lets us add components as necessary to meet the needs of a large-scale user base. For scaling SIP servers, we make use of the SRV feature in DNS⁸ and the redirect capability in SIP. A DNS SRV resource record like the one in Figure 5 lists a set of servers, ordered by priority, for each service and domain. Servers *a*, *b*, and *c* in the example are to be used first (priority 0) if possible, and server *d* is the backup server (priority 1). Among the primary servers, *a* and *b* should receive a combined total of 80 percent of the requests; server *c*, presumably the slowest of the three, is designated to handle the rest. The client uses weighted randomization to achieve this distribution.

Simple random request distribution is not sufficient because servers need to share access to the same registration information. Each server would thus have to replicate incoming registration requests to all other servers or update a common shared database. In either case, updates or lookups would quickly create a bottleneck. Given that SIP user agents typically register once an hour, the database for a wireless operator with one million phones would have to process about 280 requests per second.

We solve this problem with our two-stage scaling architecture, shown in Figure 6, which divides a domain server group into two parts: a set of stateless proxy servers and a set of clusters. The proxy servers perform stateless request routing. They might route on the basis of a hash of the user identifier; each hash-range might map to a particular second-stage server, or, to ensure reliability, a cluster of servers. The system determines the cluster member using DNS SRV. This architecture scales to any desired processing load and user population size. The system can also use SIP's redirect feature to delegate requests to servers that are less loaded.

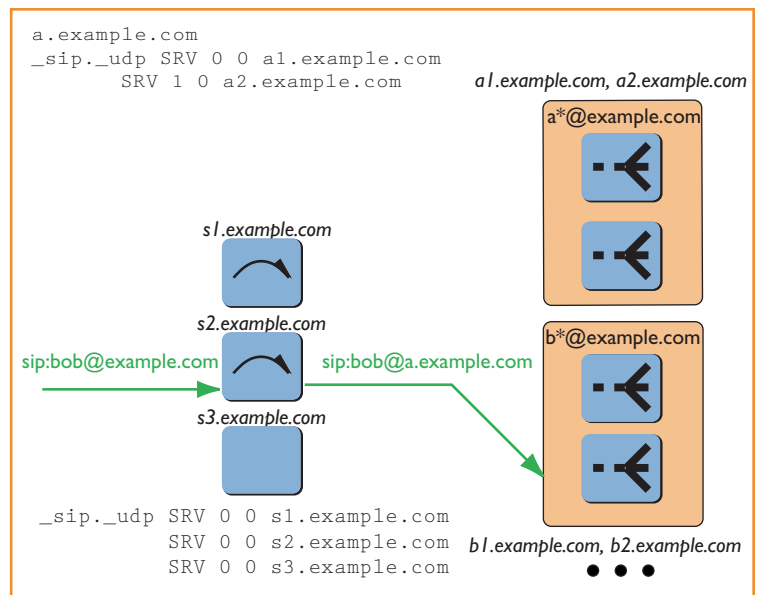


Figure 6. SIP scalability using DNS SRV. Each letter of the alphabet gets its own server cluster so that a request for bob@example.com will be routed to a server in the “b” cluster. The second-stage server then performs the actual request processing.

Our RTSP media server (rtspd) and SIP conference server (sipconf) handle media streams. Using multiple servers, each handling several tens of sessions, distributes the load and scales the system. Bandwidth might be an issue for the SIP-PSTN gateway if it handles a lot of calls, but we can eliminate this problem by using a fast 100-Mbps network interface, a low bit-rate codec (such as G.729), silence suppression,¹⁵ or a combination of these options.

Future Work

To deploy the Cinema system commercially would involve other issues related to security, billing, and QoS. One disadvantage of current PBXs is the difficulty in tailoring billing information to local needs, which usually results in having to inspect paper printouts manually. We have implemented call logging in sipd with an SQL database and Radius log files. We have also implemented an SNMP-based management agent. Cinema could provide hop-by-hop encryption and authentication using transport layer security (TLS) and secure sockets layer (SSL), but interoperating with corporate firewalls and network address translators is another challenge.

We are currently analyzing the performance of our SIP proxy and conference server. It is more difficult to evaluate the performance of SIP servers than of Web servers because finding the maximum operating rate is complicated by SIP's use of the

user datagram protocol (UDP), which causes packet loss and retransmission under overload. Also, the workload is likely to differ dramatically between registration-bound mobility service and script-processing service engines.

We plan to develop a Windows CE version of our SIP UA, making it possible to integrate wireless PDAs into the infrastructure. A short-term goal is to deploy the system through the computer science department and replace our PBX completely. We also plan to address QoS issues and provide direct SIP services for integrated access. □

Acknowledgments

Equipment and research grants from 3Com, Cisco, Clarient Communications, Lucent, SIPcomm, and Sylantro supported this work.

References

1. H. Schulzrinne et al., "RTP: A Transport Protocol for Real-Time Applications," Internet Engineering Task Force RFC 1889, Jan. 1996; available at www.rfc-editor.org/rfc/rfc1889.txt.
2. M. Handley et al., "SIP: Session Initiation Protocol," IETF RFC 2543, Mar. 1999; available at www.rfc-editor.org/rfc/rfc2543.txt.
3. *Recommendation H.323, Packet-based Multimedia Communication Systems*, Int'l Telecommunications Union, Geneva, Feb. 1998; available at www.itu.int/rec/recommendation.asp?type=folders&lang=e&parent=T-REC-H.323.
4. M. Handley and V. Jacobson, "SDP: Session Description Protocol," IETF RFC 2327, Apr. 1998; available at www.rfc-editor.org/rfc/rfc2327.txt.
5. K. Singh and H. Schulzrinne, "Unified Messaging Using SIP and RTSP," *Proc. IP Telecom Services Workshop*, AT&T Labs Research, Atlanta, Sept. 2000, pp. 31-37; available at www.research.att.com/conf/ipts2000/.
6. K. Singh, G. Nair, and H. Schulzrinne, "Centralized Conferencing Using SIP," *Proc. Internet Telephony Workshop 2001*, Columbia Univ. Press, New York, Apr. 2001, pp. 57-63; available at www.fokus.gmd.de/events/iptel2001/.
7. A. Gulbrandsen and P. Vixie, "A DNS RR for Specifying the Location of Services (DNS SRV)," IETF RFC 2052, Oct. 1996; available at www.rfc-editor.org/rfc/rfc2052.txt.
8. J. Rosenberg and H. Schulzrinne, "A Framework for Telephony Routing over IP," IETF RFC 2871, June 2000; available at www.rfc-editor.org/rfc/rfc2871.txt.
9. J. Bellamy, *Digital Telephony*, John Wiley & Sons, New York, 1991.
10. J. Lennox and H. Schulzrinne, "CPL: A Language for User Control of Internet Telephony Services," IETF Internet draft, Nov. 2001; work in progress.
11. J. Rosenberg, J. Lennox, and H. Schulzrinne, "Programming Internet Telephony Services," *IEEE Network*, vol. 13, no. 3, May/June 1999, pp. 42-49.
12. J. Lennox, H. Schulzrinne, and J. Rosenberg, "Common Gateway Interface for SIP," IETF RFC 3050, Jan. 2001; available at www.rfc-editor.org/rfc/rfc3050.txt.
13. X. Wu and H. Schulzrinne, "Using the SIP Message Method for Shared Web Browsing," IETF Internet draft, Nov. 2001; work in progress.
14. T. Richardson et al., "Virtual Network Computing," *IEEE Internet Computing*, vol. 2, no. 1, Jan./Feb. 1998, pp. 33-38.
15. W. Jiang and H. Schulzrinne, "Analysis of On-Off Patterns in VoIP and Their Effect on Voice Traffic Aggregation," *Proc. Int'l Conf. on Computer Comm. and Networking*, IEEE Press, Piscataway, New Jersey, Oct. 2000, pp. 82-87.

Wenyu Jiang is a PhD candidate at Columbia University. He received an MS in computer science from Columbia University. His research interests include real-time multimedia and voice over IP, quality-of-service measurement and management, and Internet telephony.

Jonathan Lennox is a PhD candidate at Columbia University. He received an MS in computer science from Columbia University. His research interests include user-controlled services for Internet telephony.

Sankaran Narayanan is an MS student at Columbia University. He received a BE in computer science and engineering from Regional Engineering College, Tiruchirappalli, India. His research interests include Internet multimedia and telephony, scalable distributed systems, and performance evaluation.

Henning Schulzrinne is an associate professor at Columbia University. He received a PhD in electrical engineering from the University of Massachusetts, Amherst. His research interests include Internet telephony, Internet multimedia control and transport, and performance evaluation.

Kundan Singh is a PhD candidate at Columbia University. He received an MS in computer science from Columbia University. His research interests include Internet telephony, SIP-H.323 signaling gateway, unified messaging systems, and multimedia conferencing.

Xiaotao Wu is a PhD candidate at Columbia University. He received an MS in computer science from Columbia University. His research interests include Internet telephony and communication services.

Readers can contact the authors at {wenyu; lennox; sankaran; hgs; kns10; xiaotaow}@cs.columbia.edu.