♦ Issues in Supporting Multimedia Services in SIP Networks with Mixed Endpoint Types

Maria R. G. Azada, Jim Calme, Richard P. Ejzak, and John J. MacNamara

The evolution from traditional circuit-based networks to packet-based networks is a salient feature of today's telecommunications industry. As new types of Session Initiation Protocol (SIP) networks and endpoints (e.g., 3rd Generation Partnership Project Internet Protocol Multimedia Subsystem [3GPP IMS], cable, public switched telephone network [PSTN] gateways, and 802.11-capable personal digital assistant [PDAs]) are introduced, SIP extensions are being introduced to solve the problems unique to each of them. As SIP networks mature to support mixed endpoint types, interesting issues relating to interoperability among the various endpoints arise. This paper includes a discussion of some of these issues, including out-of-band versus in-band call progress information, handling early media, support of Internet Protocol (IP) and asynchronous transfer mode (ATM) transport networks, interworking with ISDN User Part (ISUP), and support of different SIP extensions. It also explores issues of performance, reliability, and service quality that arise as networks serving mixed endpoints are deployed.

Introduction

Packet networks based on Session Initiation Protocol (SIP) and Internet Protocol (IP) are showing great promise for the telecommunications industry. Industry analysts believe this new technology will enable a large number of new revenue-generating services, reduce the cost of operating a network, and revolutionize the nature of telecommunications. As SIP networks mature to support mixed endpoint types and replace existing portions of the public switched telephone network (PSTN), interesting issues arise relating to interoperability among the various endpoints.

This paper begins with a discussion of the various endpoint types that may coexist in a SIP-based service network. It then explores some of the issues that surface when interoperating between mixed endpoint types. The issues discussed include interworking between ISDN User Part (ISUP) and SIP, call progress handling and early media, media transport, and signaling transport.

Endpoint Types

Several different types of endpoints (e.g., native SIP endpoints, gateways to PSTN endpoints, transit nodes, mobile switching centers [MSCs], and application servers) may coexist in a SIP-based network. Native SIP endpoints use SIP to set up call sessions; gateways to PSTN endpoints provide access to the

Bell Labs Technical Journal 9(3), 183–197 (2004) © 2004 Lucent Technologies Inc. Published by Wiley Periodicals, Inc. Published online in Wiley InterScience (www.interscience.wiley.com). • DOI: 10.1002/bltj.20050

PSTN, which requires interworking (I/W) to the current PSTN signaling protocols (e.g., ISUP) and bearer protocols (e.g., Pulse Code Modulation [PCM]); transit nodes provide transit communication between endpoints; and MSC endpoints provide access to wireless mobiles. The following subsections describe these endpoints in detail.

Panel 1. Abbreviations, Acronyms, and Terms

3GPP—3rd Generation Partnership Project
ACM—Address complete message
ATM—Asynchronous transfer mode
DNS—Domain name server
DTMF—Dual-tone multifrequency
FQDN—Fully qualified domain name
GUI—Graphical user interface
I/W—Interworking
IAM—Initial address message
IETF—Internet Engineering Task Force
IMS—IP Multimedia Subsystem
ISDN—Integrated services digital network
IP—Internet Protocol
ISUP—ISDN User Part
ITU-T—International Telecommunication Union, Telecommunication Standardization Sector
LCP—Local call progress
MGC—Media gateway controller
MGW—Media gateway
MIME—Multipurpose Internet Mail Extensions
MSC—Mobile switching center
PCM—Pulse Code Modulation
PLMN—Public land mobile network
PRACK—Reliable provisional response
PSTN—Public switched telephone network
RCP—Remote call progress
RTP—Real-Time Transport Protocol
SDP—Session Description Protocol
SIP—Session Initiation Protocol
SIP-T—Session Initiation Protocol for Telephones
SS7—Signaling System 7
TDM—Time division multiplexing
TrFO—Transcoder free operation
URI—Uniform resource identifier
VoIP—Voice over IP

Native SIP Endpoints

A native SIP endpoint is an endpoint that uses SIP signaling to establish a call session. SIP signaling is used to create, modify, and terminate multimedia sessions. Networks that support native SIP endpoints must make allowance for a number of characteristics common to such endpoints. For example, native SIP endpoints use a uniform resource identifier (URI) and a dynamic addressing scheme that requires domain name server (DNS) resolution. They also support a wide variety of codec types. Finally, they are not homogeneous; they may have a variety of capabilities and user interfaces that can be used to establish sessions. For example, when a callee is being alerted, a native SIP endpoint with a graphical user interface (GUI) may display a message or a picture of a phone ringing, while native SIP endpoints lacking such an interface may simply provide a local ringing tone, as most existing SIP endpoints do. Because of such differences, the progress of a call is communicated by having a terminating SIP endpoint generate a SIP signaling message and send it out-of-band to the originating SIP endpoint. This allows the originating endpoint to provide distinct behavior based on its capabilities; for example, an endpoint with a GUI may show a picture of a phone ringing instead of providing an audible ringing tone. SIP endpoints typically provide call progress indication using SIP provisional messages out-of-band.

A preconditions feature can be used to minimize ghost rings to the originator of a session. (The term "ghost rings" refers to an immediate alert that is sent to the originator of a session before the originator receives confirmation that the session is able to progress.) When the preconditions feature is used, resources are reserved before any alert or call progress information is sent to the originator of a session. Call progress handling occurs only after both the originating and terminating ends have confirmed that resources are available. In a homogeneous SIP network, the originating side must always be prepared to receive both out-of-band and in-band call progress indications. SIP endpoints may not mandate the reservation of resources when establishing a call; however, there are provisions in the protocol that make such reservation possible.

SIP endpoints assume SIP peers and IP transport; in other words, when a SIP endpoint initiates a session, it always assumes it is communicating with a SIP peer, and, when it is negotiating a call session, it always provides media connection information that assumes IP transport.

PSTN Endpoints

In contrast to native SIP endpoints, which have many capabilities, existing PSTN endpoints are very limited. The only user interface always supported by the basic residential phone is the simple twelve-button keypad. (Most phones also provide speed-dial buttons as a shortcut to keying in strings of digits.) To invoke a service, the user keys in a string of digits or special characters (e.g., # and *), each of which is represented on the keypad. The basic residential phone is only capable of signaling to the PSTN end office in-band; it does so by using dual-tone multifrequency (DTMF) tones that correspond to the twelve buttons. The end office provides all the service logic by analyzing the string of tones it receives. (The tones can be given special meanings based on the call state.) But simplicity has advantages; in the case of residential phones, the main advantage is interoperability. The customer is guaranteed that any phone purchased will work without startup problems when it is connected to the circuit network.

For more demanding users, and especially for business users, the integrated services digital network (ISDN) provides greater flexibility. While most homes in the United States still use the standard residential phone described above, ISDN phones are in common use as residential interfaces, especially in European countries. The typical ISDN phone—with its enhanced keypad and display—provides more features than the standard residential phone. For example, actions that require several keypunches on a standard residential phone can be accomplished on an ISDN phone with a single keypunch. These additional features are made possible by the use of the ISDN protocol between the user and the end office. The ISDN protocol also allows several call appearances to be active at the same time, using the same physical connection. Finally, the ISDN message-based protocol allows for the flexibility of the message set, thereby making it easy to introduce new services.

All PSTN endpoints, regardless of type, have a number of common characteristics:

- The addressing of the endpoint is fixed; the assigned E.164 number is statically associated with the physical termination (i.e., the twisted-pair) and not with the actual terminal equipment. Therefore, it is not possible for an end user to move the terminal equipment from one location to another and automatically have the phone number move to the new location. (In contrast, cellular equipment is able to provide such mobility, because the number is assigned to the terminal equipment, rather than the physical location.)
- Call progress tones and announcements are provided by the terminating exchange.
- Call progress tones and announcements must be provided to the endpoint in-band. The terminal equipment does not have the ability to generate such audible information on its own, based on the out-of-band signaling message.
- The bearer path is available shortly after a call setup is initiated.

Other Endpoint Types

In addition to native SIP endpoints and PSTN endpoints accessible via a signaling interworking node, this paper also considers SIP interworking with circuit mobile endpoints and their service nodes, which are typically MSCs in the public land mobile network (PLMN). **Table I** shows the SIP interworking attributes associated with each of these endpoints. Note that this table includes only a few of the SIP endpoint types and subtypes that have been defined to date. For example, endpoints from 3rd Generation Partnership Project IP Multimedia Subsystem (3GPP IMS) and cable voice over IP (VoIP) might also have been included in the table (and in the discussion that follows); they were excluded not because they are less important, but simply to limit the range of issues that must be considered.

Circuit mobile endpoints are typically described as PSTN endpoints, but this is only because they are adapted for PSTN interworking. When viewed as

Table I. Endpoint comparison of SIP interworking attributes.

SIP interworking attribute	Native SIP endpoint	PSTN endpoint	Circuit mobile endpoint	PLMN transit node
Service intelligence	Preferably in the endpoint	In the network	In the network	In the node
Addressing	SIP, URI, and DNS	E.164	E.164	E.164
Mobility	Inherent in registration procedure	None	Inherent in IS-41 registration procedure	None
Call progress information	Usually provided locally based on signaling	Always provided in-band within the bearer	Could provide either in-band or out-of-band	Could provide either in-band or out-of-band
Transport network	Typically IP	TDM	Packet	As needed
Media codecs	Packet codec as required	Typically G.711	Mobile packet codec	As needed

DNS—Domain name server
IP—Internet Protocol
PLMS—Public land mobile network
PSTN—Public switched telephone network
SIP—Session Initiation Protocol
TDM—Time division multiplex
URI—Uniform resource indicator

endpoints in a SIP network, they have more characteristics in common with native SIP endpoints than with PSTN endpoints; for example, circuit mobile endpoints use mobile packet codecs that can be transported natively through an IP network rather than forced into a time division multiplexing (TDM) network via transcoding to G.711. This type of operation—also known as transcoder-free operation (TrFO) [4]—makes it possible to optimize mobile networks, because it avoids unnecessary transcoding for mobile-to-mobile calls. It also makes it possible to place any necessary transcoders remotely in the network to allow for optimal transport efficiency. Mobile packet codecs also provide significant transport savings, because they require only a fraction of the bandwidth used by TDM/G.711.

Other areas for potential optimization involve the mechanisms available for handling call progress information when adapting circuit mobile endpoints to a SIP network. For example, it is not necessary for a terminating circuit mobile endpoint to provide call progress information in-band within the bearer, because SIP messaging can communicate call progress information out-of-band to the originating side, which can then provide the appropriate tones or announcements to the calling party. The same is true for any PLMN transit node. In both cases, it is more efficient to provide out-of-band indications to the originating side, which can then provide any necessary call progress information. This mode of operation makes it unnecessary to allocate bearer resources at the terminating and transit nodes and to transmit media packets through the bearer network. It also makes it unnecessary for transit nodes to allocate media resources solely to force media packets to pass through these resources on the path between the originating and terminating nodes in the bearer network. This, in turn, makes it unnecessary for the network to allocate media resources at the transit nodes and reduces the distance media packets traverse in the bearer network by allowing them to pass directly between the originating and terminating nodes.

Issues Related to Interworking Mixed Endpoint Types

Ideally, a SIP network should support all endpoint types interchangeably, but such interoperability cannot be assumed, because each endpoint type has a different way of managing such things as addressing, call progress, and media. The remainder of this paper focuses on some of the issues involved in making such interoperability possible.

ISUP-SIP Interworking

As new protocols are introduced, it becomes necessary to make them interwork with existing network protocols in order to support complete network interconnectivity. Interworking between two protocols provides a mapping of the information elements of one protocol into similar information elements of another protocol, based on a well-defined set of rules. In some cases, the mapping is straightforward, with little or no loss of information (e.g., interworking between ISDN and ISUP). However, when protocols are significantly different, the interworking becomes more difficult to accomplish without losing details in the translation.

In such cases (i.e., when the interworking does not derive all the necessary information), it is necessary to make assumptions about information and to provide default values by introducing interworking elements—which we refer to as interworking points—into the network as the call progresses between two protocols. In the case of interworking between a circuit network using ISUP signaling and a packet network using SIP signaling, the interworking point must provide the interworking at both the call-control level (i.e., ISUP-SIP) and the bearer level, in this case between TDM and Real-Time Transport Protocol (RTP). Therefore, the interworking point controls both the call-control signaling and the bearer signaling, so it must provide both the media gateway (MGW) and the media gateway controller (MGC) functions.

There are two possible interworking scenarios. The first scenario is one in which a SIP endpoint and a PSTN endpoint are communicating. In this case, only basic call services are expected. In **Figure 1**, which illustrates this scenario, the interworking point must provide a mapping between the two protocols that will allow the call to progress through the next network.

The other scenario is shown in **Figure 2**, in which a SIP packet network is used to connect, or bridge, two PSTN networks. In this case, it is expected that the services supported by the originating PSTN network will be available to the destination PSTN network, even though the call passes through the intermediate SIP network. For this reason, more than basic interworking is required in this scenario.

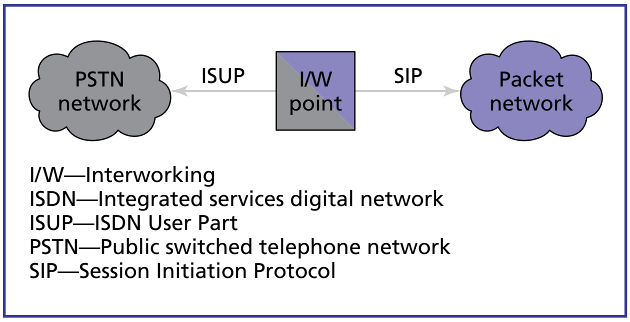

Figure 1.
PSTN interworking between PSTN and packet networks.

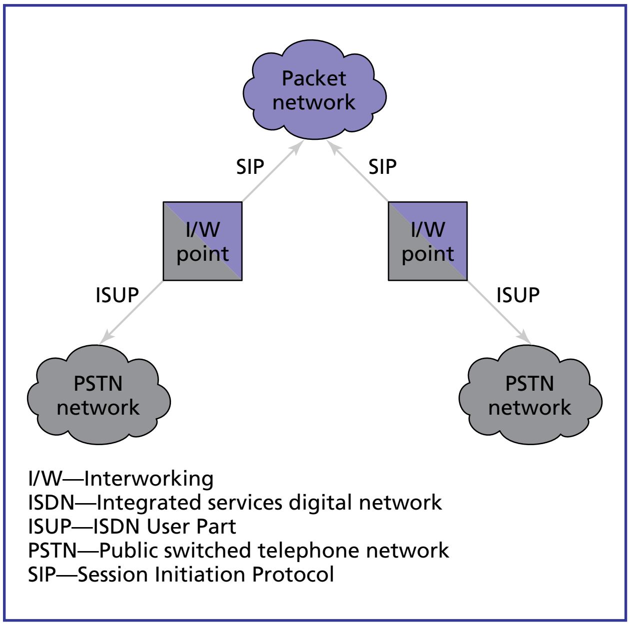

Figure 2.
PSTN interworking: bridge between two PSTN networks.

To ensure interoperability among various implementations, it is necessary to standardize the rules, assumptions, and defaults of an interworking. A number of industry-sponsored organizations are working on standardizing the interworking between SIP and

ISUP, among them the Internet Engineering Task Force (IETF), which has published RFC 3398 [2], and the International Telecommunication Union Telecommunication Standardization Sector (ITU-T), which is publishing Q.1912.5 [9]. Both documents map information elements from one protocol to the other, but the two mappings differ slightly. As is the case with SIP to ISUP interworking, only those services that are common to the two protocols can be supported across the interworking point. When a service cannot be supported, the interworking point should attempt to terminate gracefully or to deny the service.

The interworking between SIP and ISUP focuses mainly on the information needed to establish the call (i.e., the information in the ISUP IAM message and the SIP INVITE messages), because this is the majority of the information that is carried. However, a mapping of each call-related message must also be provided. **Figure 3** illustrates the interworking of ISUP and SIP information, as defined by the U.S. T1S1 ISUP-SIP interworking specification.

The *request URI* of the SIP INVITE message is expected to contain a fully qualified E.164 telephone number (i.e., a "+", followed by a country code, followed by the national number). If the number is not fully qualified, the INVITE should be rejected. Otherwise, the interworking point must make assumptions regarding the number format (e.g., in the case of private numbering plans).

Because the SIP INVITE *From* header should not be trusted as the actual identification of the calling party, an additional SIP INVITE header, the *P-asserted-ID*, was introduced in RFC 3325 [10]. The SIP network populates this header for authenticated users. When it is present, the *P-asserted-ID* header may be used to populate the ISUP IAM *Calling party number*. The SIP INVITE *Privacy* header should be used to populate the address presentation restricted indicator of

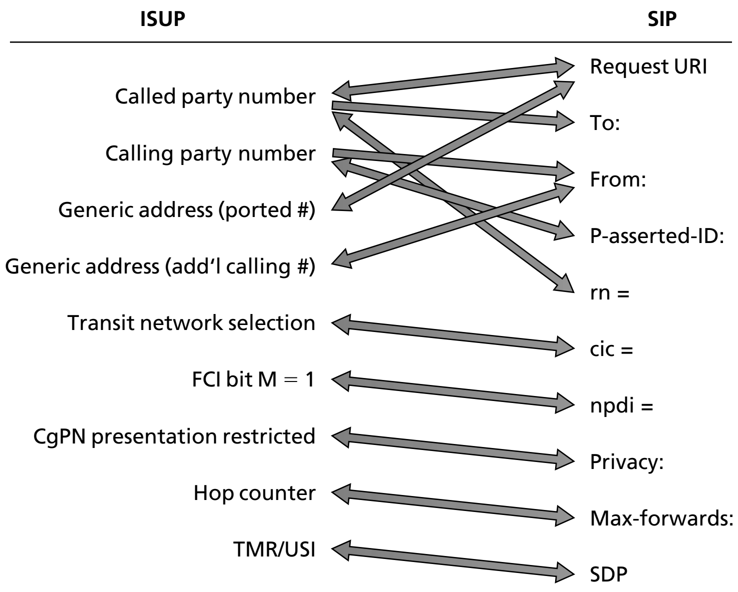

CIC—Circuit identifier code
CgPN—Calling party number
FCI—Forward call indicator
ISDN—Integrated services digital network
ISUP—ISDN User Part
NPDI—Number portability dip indicator
RN—Routing number
SDP—Session Description Protocol
SIP—Session Initiation Protocol
TMR—Transmission medium requirement
URI—Uniform resource indicator
USI—User service information

Figure 3.
ISUP/SIP parameter interworking.

the ISUP IAM *Calling party number*. This allows the privacy of the calling party to be maintained across the interworking point. The SIP INVITE *From* header may also be passed along in the ISUP IAM as an additional *Calling party number*.

Although many services are not supported because of limited interworking, there are some services—specifically carrier selection and number portability—that must be supported to meet U.S. government regulations. To meet these regulations, additional SIP extensions have been defined [5, 8]. These extensions define a carrier identification code (*cic*) to provide support for carrier selection and routing number (*rn*) and number portability dip indicator (*npdi*) [18] parameters to provide support for number portability. Both may be included on the *Request URI* line of the SIP INVITE.

While most of the mappings only involve copying information from one protocol to the other, the mapping between the ISUP IAM *Hop counter* and the SIP INVITE *Max-forwards* is a little more complicated, because the allowed ranges of the two parameters differ significantly. Therefore, the mapping between the two must scale the value being mapped proportionately, depending on the direction of the interworking.

Figure 4 shows the message mapping at an interworking point for a basic ISUP to SIP call flow. All call-control messages are interworked between the two protocols, but there are a number of service-related messages (e.g., ISUP INR/INF messages) that are not interworked. There is also no need to provide interworking for messages not related to calls (e.g., ISUP trunk-maintenance messages and SIP registration messages), because these messages are only related to the interface on which they are received.

As we have seen, only a small number of the many information elements defined by ISUP and SIP are actually supported by the interworking function. This simple mapping limits the services that can be supported when ISUP-SIP interworking occurs, which, in turn, limits the usefulness of SIP as a network protocol when interconnecting mixed networks. SIP has many advantages over the existing PSTN protocols, because it can provide complete end-to-end connectivity in a homogeneous packet network. But

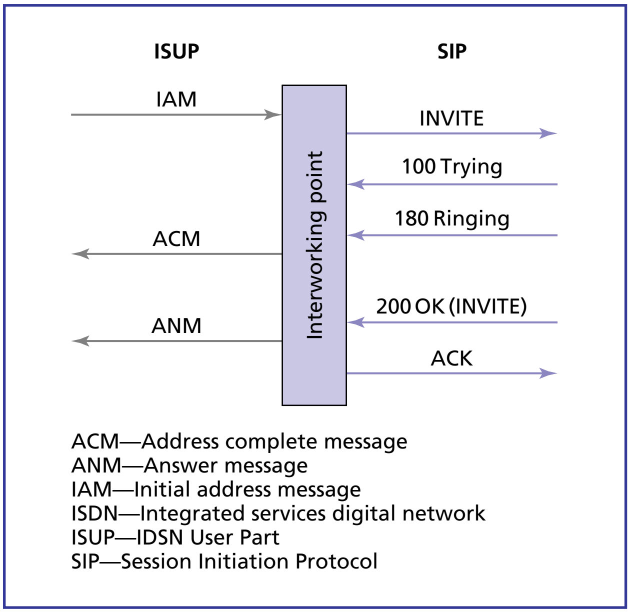

Figure 4.
ISUP to SIP protocol message interworking.

we are far from replacing the existing PSTN with a fully packetized network. Therefore, it is essential that SIP be able to provide support for more services when interworking with the PSTN. There are two possible approaches to achieving this goal:

- Continue to extend SIP, as has been done for carrier selection and number portability, or
- Provide an alternate means of transporting the ISUP information across the SIP network.

The argument against the first approach is that SIP was not intended to be a replacement of ISUP and PSTN access protocols such as ISDN. In fact, the manner in which a service is provided in SIP may be quite different from the manner in which the same service is provided in ISUP. Furthermore, simply duplicating ISUP parameters in SIP is at variance with the direction in which SIP is evolving. Nevertheless, at a minimum, information equivalent to that contained in the ISUP parameters that support those network routing and network-level services that will still be needed in a packet network must continue to be introduced into SIP.

For the reasons just mentioned, the industry has adopted the second approach and has chosen to transport the ISUP information as an attachment to the SIP message. One advantage to this approach, which

Panel 2. ISUP Message as SIP MIME Attachment within SIP Message

```
INVITE sip:+16309792000@lucent.com;user=phone SIP/2.0
          ... SIP headers ...
      Content-type:multipart/mixed; boundary=unique-boundary-1
      MIME-Version: 1.0
      --unique-boundary-1
      Content-type: application/sdp; charset=ISO-10646
            ... SDP parameters ...
      --unique-boundary-1
      Content-disposition: session; handling=required;
      Content-type: application/isup; version=ANSI00
      Content-disposition: signal; handling=optional;
           ... binary encoded ISUP message ...
      --unique-boundary-1
```

is known as ISUP encapsulation, is that an all-SIP network is not burdened with extra headers when ISUP is not involved in the call. Another advantage is that it is easy to support information (e.g., information specific to national ISUP variants or private extensions) that is not required to route the call across the SIP network.

ISUP Encapsulation

ISUP encapsulation allows ISUP-specific information to be carried transparently across a SIP network. The ISUP message is simply carried within the SIP message as a Multipurpose Internet Mail Extensions (MIME) attachment [6, 19]; it is included in the attachment in binary format. When SIP contains encapsulated ISUP, the IETF refers to it as "SIP for Telephones" (SIP-T) [17], while the ITU-T refers to it as "SIP-ISUP" (SIP-I) or profile C. When SIP does not contain encapsulated ISUP, the IETF refers to it as standard SIP, while the ITU-T refers to it either as profile A (when it is used for 3GPP) or profile B (when it is not). The interworking rules for profile A differ slightly from those for profile B.

The SIP *Content-disposition* header specifies the handling of the MIME body. When recognition of the encapsulated ISUP is critical to the call (e.g., when SIP is used to bridge two ISUP networks), the handling parameter should be set to "required." In this case, the receiving node will reject the call if it does not recognize the ISUP MIME body. When recognition of the encapsulated ISUP is not critical to the call (e.g., when the call is destined for a SIP subscriber), the handling parameter should be set to "optional." In this case, the ISUP information will be available to call-control procedures if the SIP endpoint can recognize the ISUP attachment. If it cannot, the attachment can be safely ignored without affecting the call. The IETF specifies that the *Content-disposition* handling parameter be set to "optional" for SIP-T. The ITU-T specifies that it be set to "required" for SIP-I/profileC.

The code in **Panel 2** shows how an ISUP message would be carried within a SIP message as a SIP MIME attachment, as described in RFC 3372 [16, 19].

Figure 5 shows an example call flow that uses SIP with ISUP encapsulation to bridge two ISUP networks. When an interworking point receives an ISUP message, it provides the basic ISUP to SIP interworking. It may also encapsulate the received ISUP message as a MIME body attached to the SIP message. The SIP INFO [5] message may be used to support the transport of mid-call ISUP messages along the SIP session signaling path. This message will not affect the SIP call state.

As a call progresses across the packet network, the SIP-level information may be modified, while the encapsulated ISUP may remain untouched. When

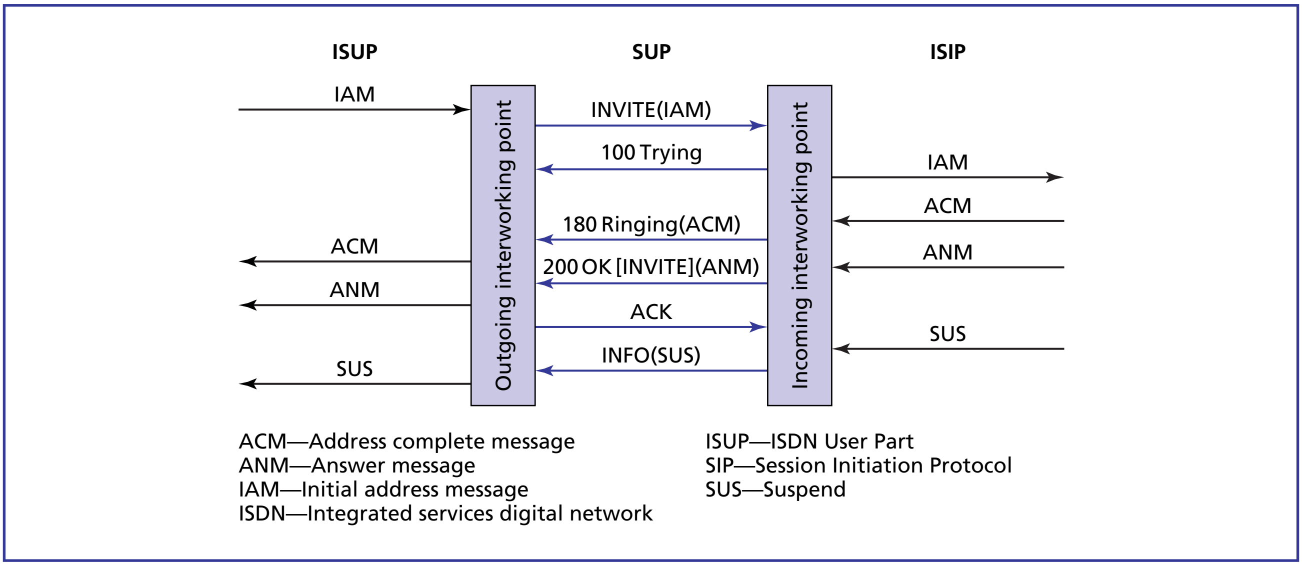

Figure 5.
ISUP encapsulation in SIP messages.

the call reaches an interworking point destination, the ISUP information is used as a template for the information to be passed to the call control function. This template is updated with the corresponding SIP-level information if that information has been modified. Thus, the SIP-level information takes precedence over any encapsulated ISUP information. If the destination is a SIP endpoint, it may or may not use the encapsulated ISUP. If the ISUP attachment is recognized, the information—possibly updated with SIP-level information—is passed on. If the ISUP attachment is not recognized, the message is processed according to SIP procedures for unrecognized MIME attachments.

Call Progress Handling and Early Media

Early media are media that are exchanged prior to the establishment of a call session. Such media may or may not contain call progress indication. The handling of call progress indication varies, depending on the endpoint type. Some endpoint types provide call progress indication locally, immediately after sending an INVITE transaction; others wait for the terminating party to trigger call progress indication. Each endpoint has predetermined methods of handling call progress indication. As we have seen, PSTN endpoints and native SIP endpoints have differing methods of handling call progress indication. Network providers that allow different endpoint types to coexist in their systems must have procedures that allow call progress indication to be handled consistently by the various endpoint types they support.

Establishing an early media path allows call progress indication to be received in band. We will refer to receiving call progress indication in band as remote call progress (RCP). Call progress may also be signaled out of band, indicating that call progress must be generated based on local policies; we will refer to this as local call progress (LCP). PSTN endpoints always assume RCP when handling call progress indication. There has never been any incentive to provide out-of-band signaling of call progress indication for PSTN endpoints, because call establishment always required the allocation of end-to-end bearer resources. But as SIP endpoints were introduced, out-of-band call progress (or LCP) became more desirable, because native SIP endpoints do not require that end-to-end bearer resources be allocated in the network in order to set up a call. Resource reservation is also not mandatory for a call to proceed, when SIP endpoints are involved. The use of LCP in a packet network makes it unnecessary to exchange call progress media packets through the packet network. Call progress indication is signaled to and generated locally by an originating endpoint, because there may not be any network bearer resources allocated.

Keeping call progress out of band effectively reduces call setup resources and the use of packet network capacity.

SIP-based networks for converged systems rely on the SIP protocol to provide interworking of the endpoints. As described in RFC 3261 [14], SIP has various methods and extensions that can be used to help accomplish this interworking. SIP 183 Session Progress Message and SIP 180 Ringing are examples of provisional responses that can be used to indicate session progress. It is important to note that SIP defines two types of responses: provisional and final. Provisional responses provide progress information for a specific request, but they are not sent reliably. Final responses provide the result of a request, and they are sent reliably. Interworking with PSTN endpoints requires that the provisional responses also be sent reliably. PSTN networks require that provisional responses be assured of reception so that the network can transition properly between PSTN call states. RFC 3262 [13] is an extension of RFC 3261 that allows provisional responses to be sent reliably. Sending reliable provisional responses is required by PSTN endpoints that always include encapsulated ISUP. It is also required for the establishment of early media to allow for the handling of call progress indication. After an INVITE transaction has been launched, it may transit multiple nodes in the network, so it could be some time before the terminating end generates a final response.

Early media indication is not easily detected using the SIP protocol, because the protocol does not have information that indicates the presence or absence of early media. There are various proposals for how a SIP user agent might detect early media. One proposal, which is mentioned in the IETF's early media draft [3], suggests an implementation that monitors the presence of incoming media at a media gateway. Such an approach works well with SIP endpoints, because they have the intelligence to perform such monitoring. However, in the case of gateways to PSTN endpoints or PLMN transit nodes, it is a less attractive option, because gateways are not always allocated when establishing calls for packet transport and, even if they are, the signaling control of the gateways is separate from the bearer, making it difficult to monitor the presence of incoming media. A more desirable implementation, as noted in [9], uses the presence of ISUP encapsulation as an indication that early media are available. The following discussion of early media handling and call progress conversion uses this approach when interworking mixed endpoint types controlled by a common SIP network. Note that all the figures referred to in the discussion assume reliable provisional responses (PRACKs), which are not shown.

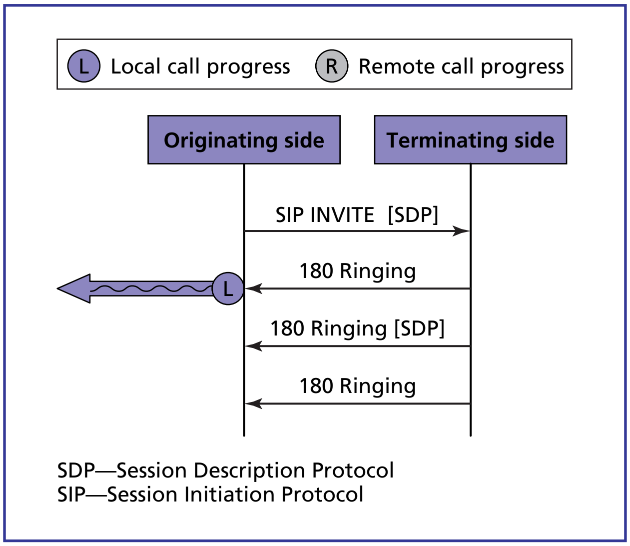

Figure 6.
Origination to a mobile endpoint (Q.1912.5 profile B).

Figure 6 illustrates the procedures at an originating side for a call or session terminating at a SIP endpoint that assumes Q.1912.5 profile B. The terminating SIP endpoint returns an out-of-band SIP signaling message (i.e., 180 Ringing) to indicate call progress. There is no ISUP encapsulation present from a SIP terminating endpoint, as indicated by profile B. This indication (i.e., 180 without ISUP encapsulation) can be used by the originating endpoint to determine that it is responsible for applying call progress handling locally upon receipt of this message. SIP endpoints do not typically generate encapsulated ISUP. Packet transport is assumed and allocation of bearer resources (e.g., a media gateway) at the terminating end is optional. We can assume that a 180 Ringing provisional response without

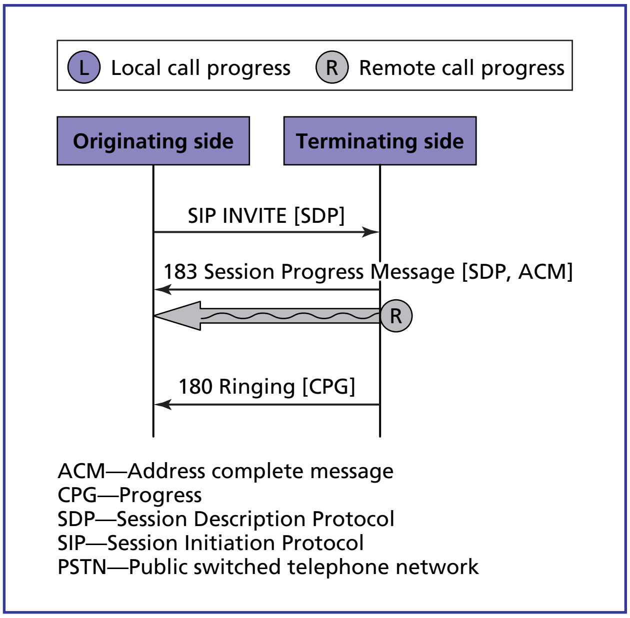

Figure 7.
Origination to a PSTN endpoint (Q.1912.5 profile C).

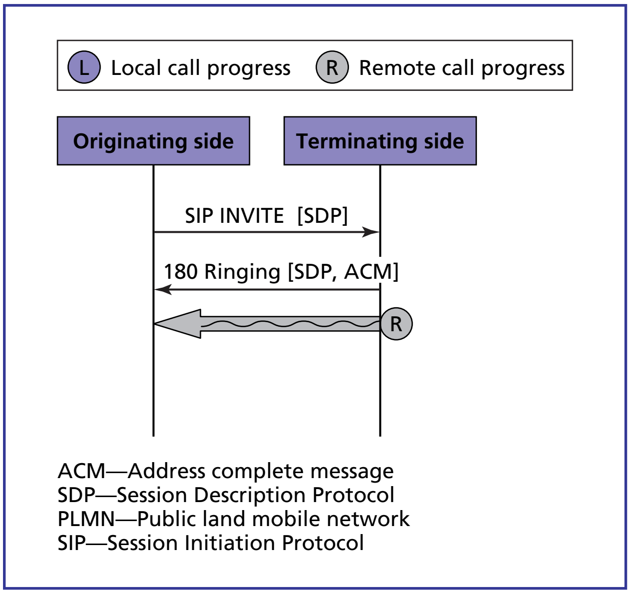

Figure 8.
Call to a SIP endpoint (profile C, handling=optional) that has traversed a PLMN transit note.

encapsulated ISUP indicates that LCP will be applied by the originating endpoint. If LCP is currently applied and another 180 Ringing provisional response is received, LCP will continue at the originating endpoint. There could be multiple stages to the call progress. In such a case, it would be possible to change the alert provided to indicate different stages of the call in progress. For example, a separate announcement might be played to the originating party to indicate that a called party is being located. Then, once the called party has been located, another announcement or tone might be played to indicate that the called party is being paged.

Figure 7 illustrates a call to a PSTN endpoint, assuming Q1912.5 profile C. An endpoint operating in profile C always sends the fully encapsulated ISUP message. The SIP 183 Session Progress Message in this figure contains an encapsulated ISUP address complete message (ACM). This serves as an indication to the originating side that RCP is available. A SIP 183 Session Progress Message also implies that call progress may be available in band. The media information in the Session Description Protocol (SDP) of a 183 Session Progress Message implies that media resources were allocated at the terminating side and in-band call progress information may be available.

Figure 8 illustrates a call to a SIP endpoint that has at some point traversed a PLMN transit node. As indicated, the endpoint is operating in profile C, because of the presence of ISUP encapsulation. This is indicated by the presence of the encapsulated ISUP ACM in the message. The presence of the encapsulated ISUP in a 180 Ringing message is used to determine that RCP may be available.

Figure 9 illustrates a scenario in which RCP is converted to LCP. We will refer to this as call progress conversion. The initial call is a call to a PSTN endpoint. (This is assumed based upon receipt of the first SIP 183 Session Progress Message.) The call is eventually forwarded to a SIP endpoint. It is assumed that procedures at the originating side support call progress information conversion. In this scenario, the originating side receives an in-band call progress indication and assumes that call progress is handled remotely. As the call progresses, a 180 Ringing message is received at the originating side. In this scenario, the originating side must be ready to toggle from RCP to LCP (e.g. by providing local alerting or ringing). In the reverse scenario (i.e., the originating side received an LCP indication and provided local ringing and then received an RCP indication), the originating side must be ready to toggle from LCP to RCP.

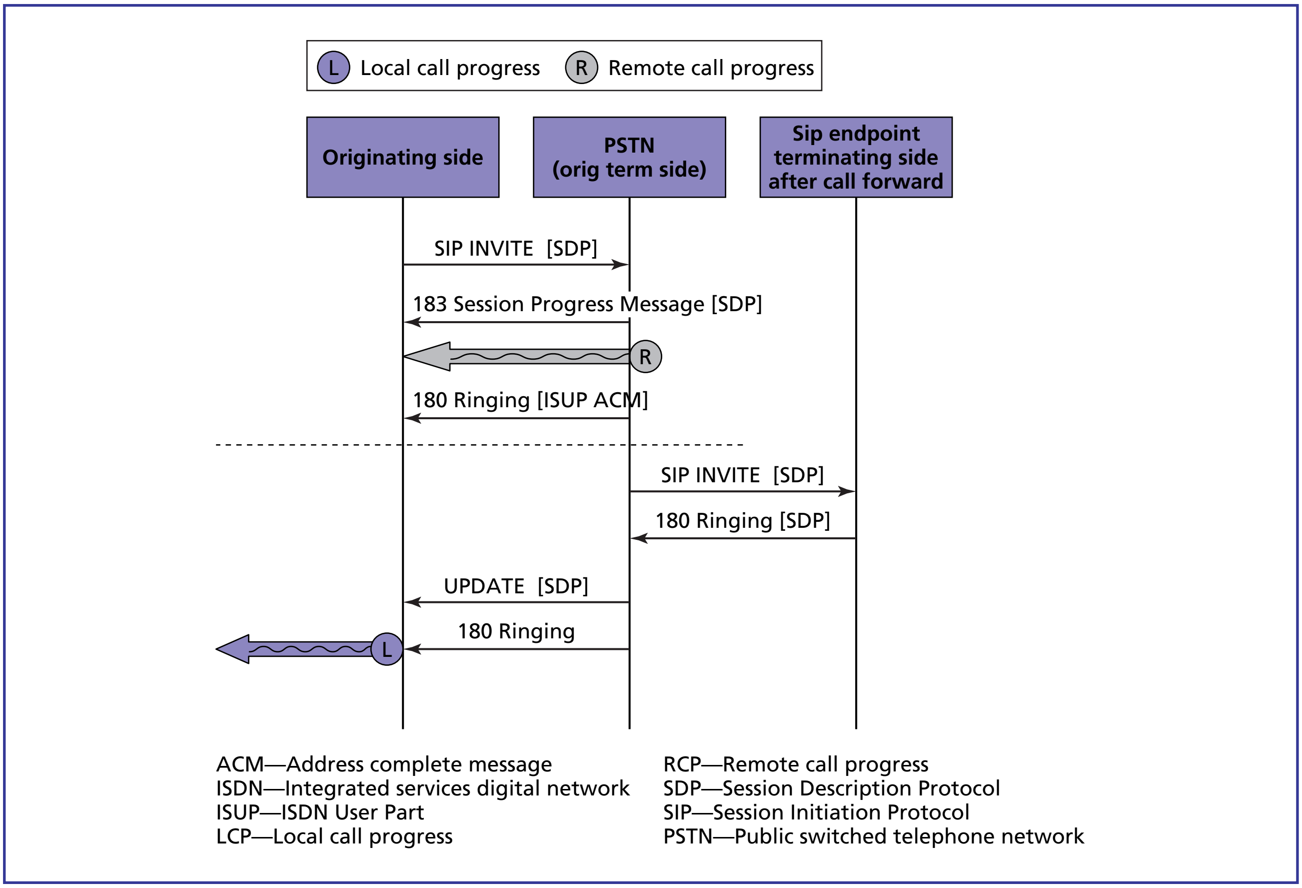

Figure 9.
Call progress conversion example (RCP to LCP): call forward from PSTN (profile C) to SIP endpoint (profile B).

Based on the standard procedures shown above, the following is a summary of call progress indication handling:

- 183 Session Progress Message with or without ISUP encapsulation indicates that call progress is available in-band and that it is handled remotely (i.e., RCP).
- 180 Ringing with ISUP encapsulation also indicates that call progress is available in-band and that it is handled remotely (i.e., RCP).
- 180 Ringing without ISUP encapsulation indicates that call progress is handled locally, based on local policies (i.e., LCP).

Media Transport Issues

Endpoint types make varying assumptions regarding transport types. As discussed earlier, SIP endpoints assume IP transport. PSTN endpoints traditionally assumed TDM or asynchronous transfer mode (ATM) transport; however, IP packet transport has also been introduced in PSTN networks, and it is becoming increasingly popular. A common SIP network that supports both ATM and IP endpoint types cannot safely assume a transport type when establishing a call.

SDP [7, 8], used as an attachment to SIP signaling, can provide a solution to the incompatible transport type assumptions of different endpoint types. RFC 3264 [12] describes how an offer/answer cycle using SDP can be used to establish a session. Negotiating media resources are presented as media offers in the SDP attached to a SIP signaling message. Transport type options are one of the many session characteristics that can be negotiated using SDP. RFC 3108 [11] discusses the use of wildcards as a means of negotiating transport type when setting up a call session.

Consider the example SDP in **Panel 3**. In this example, the connectivity to the destination is not

Panel 3. Sample of SIP Using Wildcards

```
v=0
o=HostName 28908764872 28908764872 IN
  IP4 Hostaddr
s=-
t=0 0
c= IN IP4 FQDN
m=audio $ RTP/AVP 96
c= $ $ $
a=rtpmap:96 EVRC0/8000
a= sendrecv
```

known ahead of time. The SDP includes the following information of interest:

- The session level c= line includes a fully qualified domain name (FQDN).
- The media level m= line contains a wildcard (i.e., $) in the port address.
- The media level c= line contains wildcards (i.e., $s) for the connection-level information (i.e., network type, address type, and connection address information).

This information might be provided as an SDP offer to a destination at which the transport type cannot be safely assumed. Such an offer would typically be attached to the first INVITE request. Upon receipt of the SDP offer at the terminating end, an evaluation would be performed to determine the type of connectivity that exists between the receiving end and the FQDN presented in the offer. The terminating end would then build an SDP answer based on the type of connectivity it had found. If the transport type had been determined to be IP, an IP connection would be established and the media level m= and c= lines would be returned with IP connection information (i.e., IP addresses and ports). Similarly, if other transport types had been found, the transport type connection information associated with them would be returned.

Allowing wildcards in offers has implications for media negotiation and for the minimum number of offer/answer cycles required to establish a session. The use of wildcards necessitates a second round of negotiation to establish a session. (If the endpoints support reliability of provisional responses, the second offer will typically be in either a SIP PRACK message or a SIP UPDATE message.) The first round is initiated with a wildcard offer for the connection information. This is followed by an answer indicating the far end media transport type and connection information. The originator then initiates a second-round offer to send the originating side connection information to the far end. The final answer is an acknowledgment from the far end of the second round offer.

Media negotiation procedures enforced in converged networks should be adaptable so that they can accommodate the various endpoint types the network may support. SDP can be effective in supporting mixed endpoint types with various transport type assumptions in a converged SIP network.

Signaling Transport Issues

In many cases, SIP networks are deployed or are under consideration for deployment to replace or bypass segments of the PSTN. However, the PSTN has evolved to offer a degree of reliability, a speed of call setup, a voice quality, a security, and an ease of operation that are difficult to match, while the Internet has a reputation for being unreliable and inconsistent and for providing low-quality service. Can SIP measure up to the standard set by today's PSTN?

Standards are already available that would enable SIP networks to meet or exceed the operational characteristics of the PSTN. What is missing is the operational experience needed for the technology to reach a similar level of maturity. In many respects, a mature SIP network has the potential to outperform today's PSTN. For example, signaling delay should be much less in a well-designed IP network than in an SS7. Even though signaling and bearer typically share the same transport network in SIP networks, existing methods can effectively segregate signaling traffic and can reserve sufficient network capacity to assure reliable delivery of signaling messages. SIP networks are also evolving to the same degree of reliability as the PSTN through the use of similar redundancy, operations, fault-detection, and fault-recovery techniques. Methods to secure SIP networks are also available, though they are not yet widely deployed. The primary method currently used to assure voice quality is over-engineering of the transport network. While this method can be effective, its use is only a stopgap measure until more effective quality-of-service control techniques are standardized and deployed [15].

Conclusion

SIP is inherently a very flexible protocol, and the IETF is still actively creating extensions to SIP and related protocols to support the invention of more endpoint capabilities and services. As more endpoint types take advantage of SIP for their signaling infrastructure, it becomes an ever greater challenge to deploy networks capable of interworking with all of them. This paper has shown how some of the challenges associated with the expansion of homogeneous SIP networks supporting single endpoint types to heterogeneous networks supporting multiple endpoint types can be addressed.

References

[1] G. Camarillo, W. Marshall, and J. Rosenbery, "Integration of Resource Management and Session Initiation Protocol (SIP)," IETF RFC 3312, Oct. 2002, <http://www.ietf.org/rfc/rfc3312.txt>.

[2] G. Camarillo, A. Roach, J. Peterson, and L. Ong, "Integrated Services Digital Network (ISDN) User Part (ISUP) to Session Initiation Protocol (SIP) Mapping," IETF RFC 3398, Dec. 2002, <http://www.ietf.org/rfc/rfc3398.txt>.

[3] G. Camarillo and H. Schulzrine, "Early Media and Ringing Tone Generation," IETF Internet Draft, Nov. 18, 2003, <http://www.ietf.org/internet-drafts/draft-ietf-sipping-early-media-01.txt>.

[4] B. Cyr, R. Ejzak, J. Lichter, J. MacNamara, and T. L. McRoberts, "Support for CDMA Circuit Mobiles in a SIP-Based Packet Core Network; A Step on the Path to IMS," Bell Labs Tech. J., 9:3 (2004), 217–236.

[5] S. Donavan, "SIP INFO Method," IETF RFC 2976, Oct. 2000, <http://www.ietf.org/rfc/rfc2976.txt>.

[6] N. Freed and N. Borenstein, "Multi-purpose Internet Mail Extensions (MIME) Part Two," IETF RFC 2046, Nov. 1996, <http://www.ietf.org/rfc/rfc2046.txt>.

[7] M. Handley and V. Jacobson, "(SDP) Session Description Protocol," IETF RFC 2327, Apr. 1998, <http://www.ietf.org/rfc/rfc2327.txt>.

[8] M. Handley, V. Jacobson, and C. Perkins "(SDP) Session Description Protocol - draft-ietf-mmusic-sdp-new-15.txt," IETF Internet Draft, Dec. 2003, <http://www.ietf.org/internet-drafts/draft-ietf-mmusic-sdp-new-15.txt>.

[9] "Interworking between Session Initiation Protocol (SIP) and Bearer Independent Call Control Protocol or ISDN User Part," ITU-T Q.1912.5, Nov. 25, 2003, Draft for approval.

[10] C. Jennings, J. Peterson, and M. Watson, "Private Extensions to the Session Initiation Protocol (SIP) Asserted Identity within Trusted Networks, " IETF RFC 3325, Nov. 2002, <http://www.ietf.org/rfc/rfc3325.txt>.

[11] R. Kumar, and M. Mostafa, "Conventions for the Use of the Session Description Protocol (SDP) for ATM Bearer Connections," IETF RFC 3108, May 2001, <http://www.ietf.org/rfc/rfc3108.txt>.

[12] J. Rosenberg and H. Schulzrinne, "The Offer/Answer Model with Session Initiation Protocol (SIP)," IETF RFC 3264, June 2002, <http://www.ietf.org/rfc/rfc3264.txt>.

[13] J. Rosenberg and H. Schulzrinne, "Reliability of Provisional Responses in the Session Initiation Protocol (SIP)," IETF RFC 3262, June 2002, <http://www.ietf.org/rfc/rfc3262.txt>.

[14] J. Rosenberg, H. Schulzrinne, G. Camrillo, A. Johnston, J. Peterson, R. Sparks, M. Handley, and E. Schooler, "SIP: Session Initiation Protocol," IETF RFC 3261, June 2002, <http://www.ietf.org/rfc/rfc3261.txt>.

[15] M. Siddiqui, K. Guo, S. Rangarajan, and S. Paul, "End-to-End Qos Support for SIP Sessions in CDMA2000* Networks," Bell Labs Tech. J., 9:3 (2004), 135–153.

[16] A. Vaha-Sipila, "URL for Telephone Calls," IETF RFC 2806, Apr. 2000, <http://www.ietf.org/rfc/rfc2806.txt>.

[17] A. Vemuri and J. Peterson, "Session Initiation Protocol for Telephones (SIP-T): Context and Architectures," IETF RFC 3372, Sept. 2002, <http://www.ietf.org/rfc/rfc3372.txt>.

[18] J. Yu, "Extensions to the 'tel' URL to Support Number Portability and Freephone Service," IETF Internet Draft, Jan. 2003, <http://www.ietf.org/intetnet-drafts/draft-yu-tel-url-07.txt>.

[19] E. Zimmerer, J. Peterson, A. Vemuri, L. Ong, F. Audet, M. Watson, and M. Zonoun, "Mime Media Types for ISUP and QSIG objects," IETF RFC 3204, Dec. 2001, <http://www.ietf.org/rfc/rfc3204.txt>.

(Manuscript approved May 2004)

MARIA R. G. AZADA is a member of technical staff in the Core Network Systems Engineering Department in Mobility Solutions at Lucent Technologies in Naperville, Illinois. She holds an M.S. degree in computer science from Illinois Institute of Technology in Chicago. Her current responsibilities include systems engineering, architecture, and standards support to evolve traditional circuit-based networks into packet-based networks.

JIM CALME is a member of technical staff in the Core Network Systems Engineering Department in Mobility Solutions at Lucent Technologies in Naperville, Illinois. He holds an M.S. degree in computer science from Illinois Institute of Technology in Chicago. His current responsibilities include systems engineering, architecture, and standards support for circuit- and packet-based transport and call control protocols.

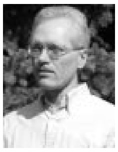

RICHARD P. EJZAK is a consulting member of technical staff in the Core Network Systems Engineering Department in Mobility Solutions at Lucent Technologies in Naperville, Illinois. He holds a B.S. degree in electronics engineering from the University of Michigan in Ann Arbor and M.S. and Ph.D. degrees, both in computer science, from Northwestern University in Evanston, Illinois. Dr. Ejzak is currently involved in product concept development in wireless IP telephony for CDMA and UMTS core networks in support of systems engineering, architecture, and standards. He holds 15 patents.

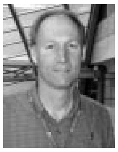

JOHN J. MACNAMARA is a technical manager in the Core Network Systems Engineering Department in Mobility Solutions at Lucent Technologies in Naperville, Illinois. He holds B.S. and M.S. degrees in computer science from North Central College in Naperville, Illinois. He is currently involved in developing product and architecture direction for circuit to packet network evolution toward the IP Multimedia Subsystem. ◆