



ELSEVIER

Computer Networks 000 (2000) 000-000

COMPUTER
NETWORKS

www.elsevier.com/locate/comnet

CitiTime: a system for rapid creation of portable next-generation telephony services [☆]

Farooq Anjum, Francesco Caruso, Ravi Jain ^{*}, Paolo Missier, Adalberto Zordan ¹

Applied Research, Telcordia Technologies, 445 South Street, Morristown, NJ 07960, USA

Abstract

We present the architecture, design and experimental research prototype implementation of *CitiTime*, an open system architecture for the rapid development of advanced next-generation telephony services that overcomes some of the limitations of the current closed PSTN architecture and service model. *CitiTime* allows communication sessions to be set up over the PSTN, the Internet, or a combination of both. Services can be provided by multiple cooperating distributed service providers, some of whom may use third-party *software components* which can be “plugged in” or even dynamically downloaded from the network as needed. This allows advanced services to be deployed and delivered to users rapidly, a crucial requirement in the increasingly competitive telecommunications services marketplace. *CitiTime* is built upon an object-oriented call model called *citi call control (CCC)* which we have defined as a small set of extensions to the standard Java telephony API (JTAPI) call model. JTAPI is designed primarily for centralized, single provider, call center type applications. Our extensions provide support for multiple, distributed providers as well as advanced services. CCC hides details of underlying call-state management, protocols and hardware from applications. The *CitiTime* prototype software is currently operational in our laboratory. We briefly describe its current implementation as well as future work to address issues such as fault tolerance. © 2000 Elsevier Science B.V. All rights reserved.

1. Introduction

It is increasingly being recognized that the current public switched telephone network (PSTN) infrastructure has an architecture that is outmoded in several respects [1]. In this paper, we present the

architecture, design and prototype implementation of *CitiTime* ², a system for the rapid development of advanced next-generation telephony services that overcomes some of the limitations of the current PSTN architecture and service model. In particular, *CitiTime* allows communications sessions (voice, video, multimedia, etc.) to be set up over the PSTN, the Internet, or a combination of both. More importantly, services can be provided by multiple cooperating distributed service providers, some of whom may use third-party *soft-*

[☆] A preliminary version of this paper was presented at the IEEE Conference on Open Architectures and Network Programming (OPENARCH), New York, 26–27 March, 1999.

^{*} Corresponding author. Tel.: +1-973-829-3178; fax: +1-973-829-2645.

E-mail addresses: rjain@telcordia.com (R. Jain), zordan@li-sa.record.unipd.it (A. Zordan).

¹ This work was done while the author was a summer intern at Telcordia Technologies.

² *CitiTime* (pronounced city-time): *CitiTime* is not an acronym; although it could stand for Computer, Internet and Telephony Integration (the Time for which is ripe).

ware components which can be plugged in or even dynamically downloaded from the network as needed. This allows advanced services to be deployed and delivered to users rapidly, a crucial requirement in the increasingly competitive telecommunication services marketplace.

Our contributions in this paper can be summarized as follows. We argue that next-generation telecommunications systems should allow rapid creation of advanced, portable services by third parties using interchangeable software components. The CitiTime system represents a step towards this goal. Intelligence is assumed to be widely dispersed and available but not completely ubiquitous; if the user does not have access to a smart terminal, services can be provided, albeit with some restrictions, by means of proxies located in the network. The software architecture of CitiTime allows services to be deployed by means of *software components* that can be developed by third-party software vendors and added to the set of available services transparently, while a session is in progress.

In order to provide these facilities, CitiTime needs a flexible and extensible *call model*, a specialized virtual machine for the development of telephony-related applications. To this end, CitiTime itself is built upon an object-oriented call model called *citi call control (CCC)* which we have defined as a small set of extensions to the standard Java telephony API (JTAPI) call model [2]. While JTAPI is generally intended for use in a single-platform environment (e.g., a PBX or a call center), CCC is designed for use in a distributed environment, and can thus be regarded as “distributed JTAPI”. CCC also hides the underlying communication protocols and architectures, providing a high-level object model for maintaining and managing the state of a call session. In particular, CCC allows different underlying session initiation and control protocols to be used, e.g., H.323 or SIP [3]. In our prototype, we have designed a version of SIP, called *extended SIP*, for session initiation which defines a small set of extensions to SIP for supporting advanced services. The CitiTime prototype implementation is currently being used in our laboratory for voice, text, whiteboard, and image-exchange sessions.

Our team is actively involved in the effort to define the JAIN call control (JCC) API, for which Telcordia is the expert group lead. The JCC API as well as the call control API defined by Parlay have many similarities to JTAPI, including many of the same objects and methods as well as Java programming patterns. For a description of the JCC API, as well as its relationship to the Parlay APIs, see [4,5]. The CitiTime call control (CCC) API described in this paper is similar to JCC but differs in many details. Firstly, it is much closer to JTAPI, as it only makes modifications to the connection object finite-state machine. Secondly, CCC is explicitly designed and implemented with the notion of being able to execute at smart end-devices, while JCC is intended to execute on call agents in the network core or periphery. Nonetheless our experience with the CitiTime prototype has significantly influenced our efforts in the JCC expert group.

In practice, JAIN and Parlay today provide the most promising migration path from IN systems for enabling third-party services. Both the JAIN and Parlay APIs incorporate many aspects of JTAPI in their approach, just as is done in CitiTime.

The paper is organized as follows. In Section 2, we present the background and motivation for this work in more detail. In Section 3, we discuss the main entities in the CitiTime architecture from a services point of view, and how services are deployed and implemented. The CCC call model and the use of the SIP protocol are described briefly in Section 4, using dynamic service download as an example. We also describe how to provide a different type of advanced service, namely the exchange of electronic business cards between users in the midst of a call, which requires multiple parallel sessions to be set up between users. In Section 5, we provide some more details on the implementation, using the unified modeling language (UML) [6], to describe the associations and interactions among the objects in the system. Section 6 briefly discusses the design as well as directions for future work. In Section 7, we briefly discuss how CitiTime is related to industry efforts like JAIN and Parlay for enabling third-party

service creation in converged networks, and in Section 8, we end with brief concluding remarks.

2. Background and motivation

The PSTN implicitly assumes that the *computational resources*, the *service creation process*, and the underlying *business model* for telecommunications services are centralized. In terms of computational resources, the current PSTN assumes that the resources available at the user terminals, or “customer premises equipment” (CPE), are extremely limited, so that all the intelligence required to provide services is centralized in network switches, databases, and operations support systems. In terms of the *service creation process*, the intelligent network (IN) architecture represented an important advance when it was introduced in the PSTN. IN separated service development from switching, allowing service logic to be developed more quickly and placed in specialized network databases (e.g., the intelligent service control point, or ISCP) while switches could be optimized for speed and efficiency. However, IN services are typically created on specialized platforms, using specialized languages and programming models, by specialized personnel employed by a telecommunications service provider. Hand-in-hand with this assumption was that of a *business model* in which a centralized service provider, one per geographical region, controlled the available communication resources and provided service.

Clearly, advances in hardware and software technology, the rise of the Internet, and market deregulation have combined to make the assumptions underlying the PSTN invalid. One of the likely scenarios for next-generation telephony is one that blends feature from the old PSTN world with those from the growing domain of IP-based data networks. Many differences exist, of course, between the two networks, but we are concerned specifically with differences at the service level. The PSTN architecture is *network-centric* and assumes that the service logic, or “intelligence”, resides within the network itself, normally in the form of programmable network elements and of centralized systems, such as the

ISCP, which implement user-level services. IP-based networks, on the other hand, are *application-centric*, in that the network passively provides connectivity among hosts, and is largely unaware of the applications that run on those hosts.

Currently, market forces are increasingly calling for these network models to converge, with consequences on the way user services are defined, developed and deployed. In a broad sense, the two approaches are complementary and are already converging: the PSTN, designed to provide primarily basic telephony, has been evolving to support new user-level, “intelligent” services. The Internet, traditionally used for client-server and peer-to-peer data applications, is increasingly being used for telephony and telecommunications services. One of the expected effects of this convergence is that the service logic will now be spread throughout the network. This means that rapid development and deployment of telephony services becomes possible, and that network providers and the service providers no longer need to coincide. One illustration of rapid deployment of an advanced service facilitated by CitiTime is the following:

Example scenario: dynamic service download. In this scenario, a service can be activated for one-time use on a customer’s terminal. If the terminal is not set up for the service, the necessary software can be downloaded and installed just prior to use. The main steps in the proposed scenario are as follows: (1) Alice, the user of terminal *A*, invites Bob, who is using terminal *B*, to a communication session where they can use a whiteboard to share information; (2) *A* issues the whiteboard session invitation to *B*; (3) unfortunately, *B* does not have the whiteboard software installed, and it informs *A* accordingly; (4) *A* suggests a location on the Internet where the software can be obtained; (5) *B* downloads the software, negotiating one-time use payment with the software provider if the software is not public-domain; (6) *B* sets up a whiteboard session with *A* so that Alice and Bob can communicate. Note that Alice and Bob have largely been spared the details of this session negotiation and dynamic service download.

Several variations of this scenario are possible. For instance, Bob may configure his terminal to

consult him before paying for one-time service activation fees if the fees exceed some threshold. Alice may authorize her own terminal to download the whiteboard software to Bob's terminal. Another possible situation is as follows. A third party, Carol, is interested in having *A* and *B* establish the session. For instance, Carol may be Alice and Bob's common supervisor. In this case, Carol may take an active role in providing Bob's terminal with information about suitable providers that can supply the services required for the session.

Fig. 1 illustrates the next-generation telephony scenario we assume. The PSTN and the Internet will continue to co-exist, and will interface with each other through gateways. On the resulting combined network, various types of terminal equipment will exist, ranging from "black" phones to personal workstations. A user will be able to enjoy a variety of advanced services, some of which are offered by multiple cooperating providers that may or may not be also be PSTN providers or Internet service providers (ISPs). In some cases, users or the service providers themselves will adopt software developed by third-party independent software vendors (ISVs), and from time-to-time providers may download services from other providers on behalf of their users. A user who is using a black phone will be able to obtain some services by means of *proxies* connected to the network, where the proxies themselves may be offered by service providers. For instance, while black phones

may be unable to support video-based services, some of the more advanced audio services can still be available to them through the proxy.

In this context, we have developed a software architecture and system, called CitiTime, that allows for the rapid creation of advanced, portable telephony applications by third-party providers. We give an overview of the CitiTime system in the following sections.

3. CitiTime service architecture

In this section, we describe how services are deployed and activated in the CitiTime architecture; in the following section, we will discuss the CCC call model used to implement service deployment and activation.

The CitiTime architecture consists of two main software entities: *endpoints* and *providers*. An Endpoint has a single unique logical address (which may be a phone number or an Internet address) in the system, and consists of one or more user terminals, a terminal manager (TM), and possibly a proxy server. A user's terminal may be a black phone, a wireless handset, a laptop, or a workstation. Because the terminal itself may not be able to support the processing logic required by the TM (this would be the case for a black phone), in our design, the terminals and their TMs are distinct but associated entities. Namely, the TM is software that either runs on the user's terminal if

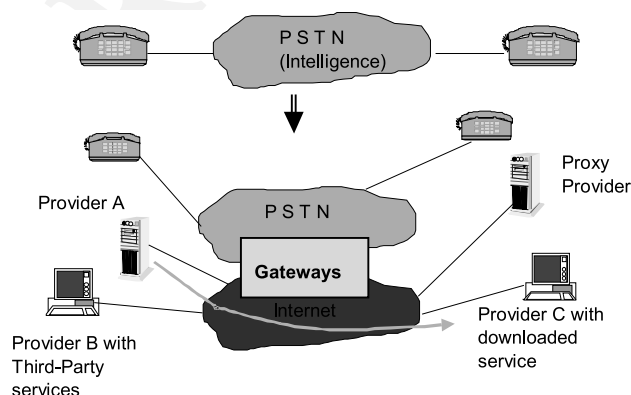


Fig. 1. Evolution of the telecommunication system.

the terminal supports such functionality, or on a proxy server in the network. Thus, from the point of view of the provider, each terminal can be uniformly addressed through the associated TM, regardless of its actual capabilities. Additionally, this decoupling allows for better scalability and ease of distribution of the overall architecture. For ease of exposition in this paper, we will assume that an endpoint consists of a TM running on a single user terminal, although other configurations are possible, and refer to endpoints and TMs interchangeably.

The CitiTime architecture is based on a network of distributed, interconnected providers (e.g., B and C in Fig. 1). While endpoints have a single logical address, providers manage disjoint sets of such addresses, called domains, and provide network connectivity and session management for the corresponding endpoints. Fig. 2 provides an overview of the relationship between endpoints, providers, and plug-in services. The CCC provider manages the call processing logic, defined basically in terms of a finite state machine, among peer TMs. TMs whose logical addresses belong to the provider's domain are shown attached to the provider.

Services can be hosted by providers or directly by the TMs, or both. Thus, like services traditionally provided in the PSTN, a service can be permanently installed and constantly active at a provider. For instance, a translation service for toll-free numbers (e.g., numbers with an 800 area code), which maps 800-numbers to a physical access line number (or an Internet address), can be

offered by a provider for its domain. The number translation can take place in a transparent way during a normal session setup when the caller application requests a connection to a peer identified by the 800 number. (The actual 800-number database could be installed at the provider, or the provider could itself access a database offered by a service provider via the network.) Alternatively, a service can be provided by the user's TM. For instance, a speed-dialing service, in which important or frequently dialed numbers are mapped to single digits (represented as hardware or software buttons on the user's terminal) could be provided by the TM. Whether a particular service is provided by a TM or a provider depends upon several factors. One such factor is that a service that is likely to be shared by many TMs (e.g., 800-number translation) may be better located at the provider, while services that are very specialized or user-specific may be better located at a TM.

In CitiTime, service and even activation can take place as part of the session setup. An instance of a service can be downloaded on demand, and activated only for the duration of a session, as in the whiteboard service of the dynamic service download scenario. The whiteboard application can be a short-lived service obtained for the duration of that session. Its lifetime and the provider designated to supply the service can be specified during session setup.

As shown in Fig. 2, services are generally software components that plug-in to the TM or the provider. Although the idea of using plug-ins borrows in part from the familiar notion of web

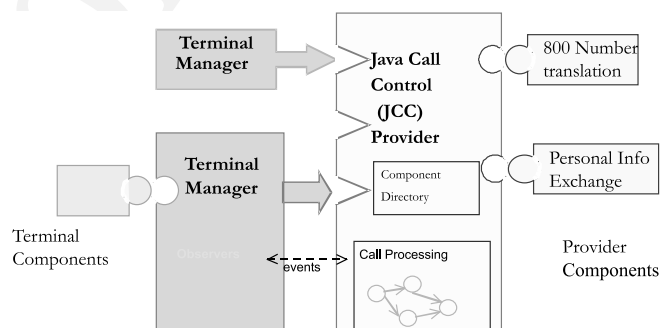


Fig. 2. CitiTime service architecture.

server plug-in, what distinguishes our model is the active management of the attached third-party software, namely activation, deactivation and access control based possibly on a subscription model. Two types of dynamic resources provide access to services. The first consists of terminal components that can be attached to a TM statically, or dynamically during call setup, and typically are user-specific or specialized services. The second type of resource, provider components, implement services that are normally shared among TMs in the domain. A terminal component is implemented as a software component (e.g., a *JavaBean* [7]) that is analogous to (or, depending upon the service, could in fact actually be) a web applet [8]. A provider component is implemented as a software component (e.g., a *JavaBean* or a “CORBA Bean” [8,9]) that is analogous to a web servlet [10] (one simplistic definition of servlets is that servlets are to web servers what applets are to web browsers). An example of a terminal component is a *JavaBean* that provides a speed-dialing service. An example of a provider component is a *JavaBean* that maintains a cache of 800-number translations, and if a dialed number is not available in the cache, contacts a database offered by an 800-number service provider to obtain the translation.

The motivation for implementing services as software components is that, in principle, different third-party software vendors can provide competing implementations of the same service. For instance, a provider may remove one vendor’s implementation of 800-number translation, and replace it with another vendor’s implementation that offers better functionality (e.g., is more up-to-date), performance (e.g., better caching or search algorithms) or cost. Current software component technology like *JavaBeans*, and enterprise *JavaBeans* [8,11], offers support for discovering the attributes of a software component at run-time by means of standardized naming conventions, and given further standardization of the service interface, allowing components to be “plugged-in” for use dynamically.

The issue of managing software components in CitiTime is similar to that of managing applets and servlets in a web-based client–server environment.

In general, CitiTime services are registered as software components in a component directory.

In the case of terminal components, the current CitiTime implementation only allows for services that set up communication sessions with different types of media. Thus, services are represented by media types, which are described using the familiar notion of MIME types. For instance, suppose that during call setup the caller requests the use of a whiteboard. A specific MIME type, say “media/whiteboard”, is used for the request. On the called party’s side, if a terminal component is associated with the requested type, then the call can be immediately accepted and the component can be activated. Otherwise, an attempt will be made to obtain a resource that can be associated with that MIME type, using dynamic service download. The notion of MIME types has so far proven to be sufficiently general to encompass all the media types we have experimented with in our prototype. Notice that in our model, MIME types are only used to qualify a request for service activation, for instance, with the INVITE message and to negotiate a common media type, but they do not by themselves trigger service activation. Thus, the type is simply an attribute of a SIP request. The “trigger points” in this prototype are currently limited to INVITE messages.

Provider components are registered in a component directory and can be activated either by a TM or implicitly by the provider. The 800 number translation service is an example of the latter. A request for a connection to an 800 number triggers the service invocation on the provider. Once the service invocation returns with the translated destination number (DN), the provider completes the connection. Notice that there is nothing preventing a provider implementation from exposing the service to the TM for a direct invocation. We omit further details of the use and maintenance of the component directory due to lack of space.

The third-party component model described raises some issues about service and feature interaction, that is, the possibility that two services that are each well-defined and correctly deployed may interact in undesirable ways when they are activated for the same user. While this service architecture does not make any specific provision for

dealing with the problem, one can argue that interaction management can be simplified by executing the external services on a common service platform controlled by a single service manager.

4. CitiTime call model

The call model used to implement CitiTime is called CCC. By hiding the details of the underlying communication protocols and architectures, CCC provides a suitable abstraction of the telephony network (the combined PSTN/IP network), and offers a high-level object model for maintaining and managing the state of a call session. Specifically, CCC builds upon complementary ideas developed by the computer telephony integration (CTI) and Internet telephony (IT) communities.

The CTI approach is oriented towards developing portable software for applications such as call centers, PBXs, etc. For example, Sun’s Java telephony API (JTAPI) [2] provides applications with a standard call model for maintaining call state, and hides the hardware API or other API (e.g., TAPI, TSAPI, etc.) from the application (see the right side of Fig. 3). In contrast, the IT approach (left side of Fig. 3) is oriented towards developing protocols (e.g., SIP, H.323) that allow interoperability and communication between software running on user terminals or gateways.

JTAPI itself, while offering a convenient abstraction for thinking about next generation telephony and networking, is somewhat limited. In particular, JTAPI seems to be oriented towards

providing support for developing applications in two types of scenarios. The first scenario is where applications run on a single platform (e.g., a PBX). The second scenario is where applications run on a platform that is “horizontally partitioned”, i.e., the higher layers of software (the application and the JTAPI layer) communicate via Java remote method invocation (RMI) with the lower layers (e.g., TAPI and the hardware) over a network. In both cases, a JTAPI provider is assumed to be in control of all the legs of a call (they all hang off the same call object managed by the provider). While this assumption may add to the convenience of managing a centralized call center, it is not realistic in the broader setting of IT. In particular, there is no provision in JTAPI for sessions that involve more than one provider, and for provider-to-provider communication for the purpose of session management.

CCC is a “distributed JTAPI” that provides the communication support for multiple providers to coordinate and interact in order to provide advanced services. For example, the JTAPI model assumes that the provider object does not change throughout the lifetime of a call. It is not clear how the JTAPI model could be used if the user wants to have different providers for different portions of the network traversed by a single call or different legs of a multiparty call. CCC overcomes these limitations.

The main goals of CCC are as follows. Firstly, CitiTime aims to provide greater flexibility in service creation, e.g., support dynamic service download, one-time service activation, and multi-

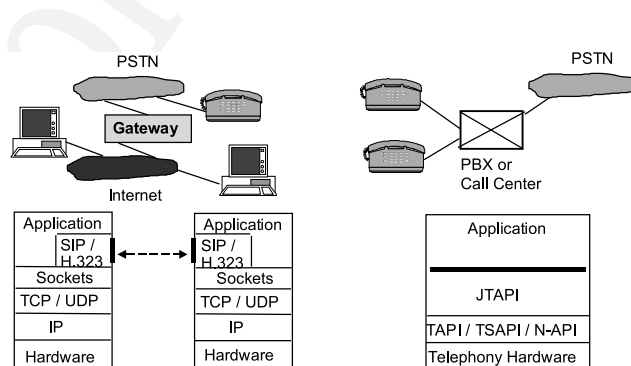


Fig. 3. IT vs CTI approaches.

ple parallel sessions between two users. Secondly, it aims to facilitate rapid development of applications by providing higher layer abstractions for call processing (rather than simply sockets, network protocols or JTAPI). Finally, CitiTime aims to provide application portability across various communication protocols, e.g., SIP or H.323, and across platforms, using Java.

The CCC architecture is organized by layers of abstraction. At the top layer, CCC offers a high-level API to applications (see Fig. 4), thus freeing them from the tasks of call-state maintenance. An intermediate protocol abstraction layer decouples the implementation of the API primitives from the underlying session control protocol. Thus, although we use SIP, the applications are independent of our choice, and the protocol can change without any major impact on the upper layers. Observe that CCC also essentially provides a

means to integrate the IT and CTI paradigms. For instance, as illustrated in Fig. 4, CCC can be implemented at a smart user terminal or proxy, where it can be used for advanced IT applications. In addition, it could be implemented at a gateway that interfaces between the PSTN and the Internet, using telephony hardware and APIs for access to the PSTN and IP-based protocols and data communications hardware for access to the Internet.

In the rest of this section, we define the CCC model in more detail. We first briefly describe the JTAPI core model and then introduce the CCC extensions.

4.1. JTAPI core model

The JTAPI [2], is a portable, object-oriented interface for Java-based computer-telephony applications. The API defines a core call model to support basic call setup, and a number of extensions, mostly designed to model call center features, multiparty conference calls, call routing, etc. The core model consists of a few telephony classes and their relationships, as shown in Fig. 5. Each object in the figure corresponds to a physical or a logical entity in the telephone world. The provider is an abstraction of a telephony service provider. A provider class manages call objects, representing calls at various stage of progress.

Terminal objects represent the physical endpoint of a call, while address objects are logical endpoints. Notice that each address can be associated with multiple terminals and vice versa, reflecting the standard configuration for a call

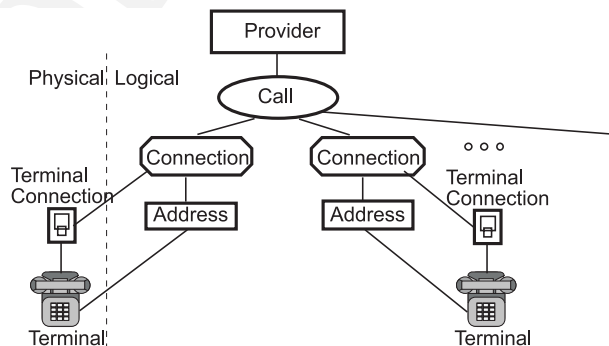
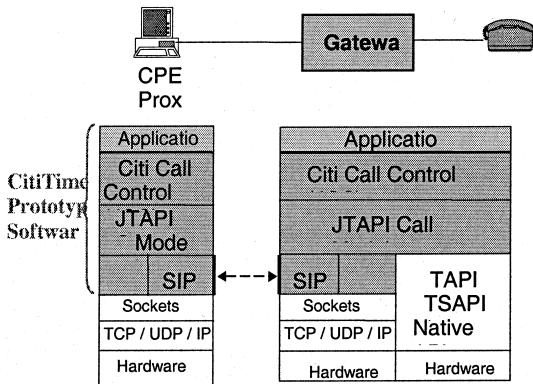


Fig. 5. Objects in the JTAPI core call model.

center. Connections model the logical relationship between a call and an address. For a multiparty call, one connection object is associated to each leg of the call. Finally, a terminal connection represents the relationship between one connection and one physical terminal.

The state of a telephone call is maintained by finite state machines associated with call, connection and terminal connection objects (*oTC*) (e.g., when a call is answered by the called party, the originating connection moves to the CONNECTED state). The complete definition of the state machines is part of the published JTAPI specifications.

4.2. Call model extensions

A TM in CitiTime is a user application that corresponds to a single JTAPI address, plus all the terminals that map to that address. Thus, a TM is addressed using a JTAPI Address, and a terminal becomes one of the resources available to the TM. Each terminal connection models the state of that resource.

The state machine for the JTAPI connection object has been extended in order to support media negotiation during call setup. The number of API method calls available to the calling parties for conducting the negotiation has likewise been

extended (a complete description of those is omitted for reasons of space). Fig. 6 illustrates the new state machine. The new states introduced are *Sender_Consult*, *Sender_Suspend*, *Receiver_Consult*, and *Receiver_Suspend*. All state transitions occur in response to events that are generated by the parties involved. One instance of the state machine is instantiated for each connection (leg) of a call. The sender and receiver states are used by the caller’s and receiver’s connections, respectively.

4.2.1. Dynamic service download

We illustrate the extended connection state machine as well as our extensions to SIP using the dynamic service download example (Fig. 7).

The caller initially sends an INVITE message to a party specifying the use of a particular media type. If the callee does not have a resource corresponding to that media, it can respond with an UNSUPPORTED_MEDIA message, which moves the caller’s connection state to *Sender_consult*. In this state, the caller’s TM has a chance to provide a suggestion to the callee, indicating a service provider, reachable through a URL, which may be able to supply the resource. If the callee accepts the suggestion, it initiates a negotiation with the provider. Since this operation takes place during call setup but is in fact conducted with a party not involved in the call, the connection for the callee

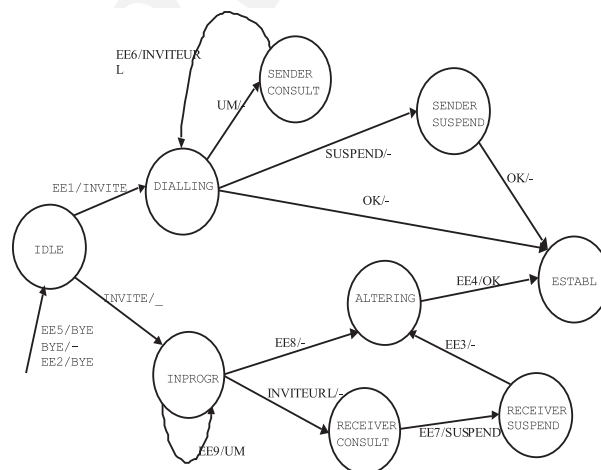


Fig. 6. State machine for the extended call model. EE1: caller initiates call attempt; EE2: time out; EE3: caller concludes action (e.g., download); EE4: callee accepts call; EE5: user hangs up; EE6: suggest action (e.g., URL for download) to callee; EE7: callee authorization; EE8: callee supports media.

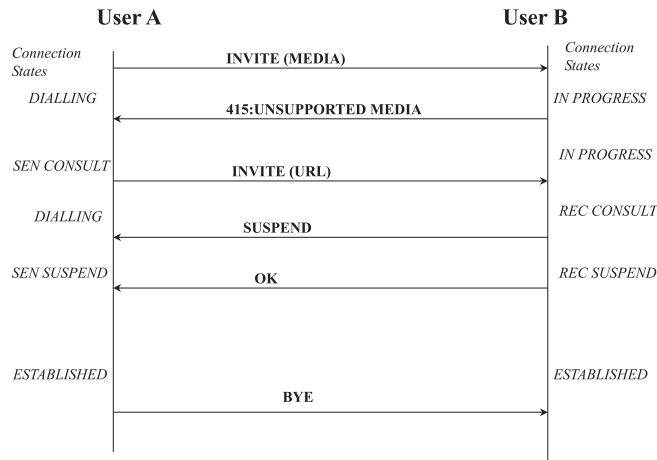


Fig. 7. Message flow for dynamic service download.

moves to the *Receiver_Consult* state. In this state, the human user at the callee end of the call may be consulted, e.g., to authorize payment for or download of the service. The caller is notified of the suspension through a *SUSPEND* event, and its connection is moved to *Sender_Suspend*. When the callee has terminated the acquisition of the resource, it notifies the caller. In case of an *OK* message, the caller then resumes its processing by moving to the standard (JTAPI) *established* state.

Notice that, at any point of time, either party can unilaterally decide to disconnect the call and return to the *IDLE* state. Also, various timeouts, not indicated in the figure, are set up in order to provide an upper bound on the negotiation time.

The remaining standard state transitions are defined in the JTAPI documentation [2].

4.2.2. Multiple concurrent sessions

A second example of the use of CCC for programming complex user-level sessions, together with media negotiation, is the support for concurrent sessions between the same TMs. Consider the following example scenario of *business card exchange*. Suppose that, during a session (after call setup), two users want to exchange their business cards from within that session. This is, incidentally, a common problem when business is conducted over the phone and a buyer is repeatedly requested to supply her name, phone number,

address, etc. (Notice that a business card information format is in fact being standardized [12], and software for performing this function in limited domains, e.g., between pairs of PalmPilot hand-held personal computers, is available.) Suppose also that such a card exchange service is in fact available to the users through a third party.

One way to accommodate this scenario in *CitiTime* is to consider business cards as a new media type. Supporting mid-call negotiation to exchange this media, however, would further complicate the call model. The alternative we use is to start a new call, in which the new media is included during setup, and then manage the logical relationship between the old and the new call. This alternative provides an opportunity to introduce and experiment with dependent relationships between calls. It also simplifies the design of the provider for the support of complex interactions, and offers more flexibility to the users. The same address and terminals share a number of active calls. Each call connection with the endpoint is represented by a different connection and terminal connection object.

Notice that the implementation of multiple sessions can be very resource-intensive. However, from the design standpoint, we believe that object-oriented principles make the implementation much simpler and more intuitive than it would otherwise be under the IN model.

A higher level session between TMs consists of all the calls shared by those TMs. Each call can be handled independently, or an explicit parent–child relationship can be enforced among them. In the case of the business card exchange, for instance, suppose that one or more sessions are setup to carry out subordinate tasks under the scope of the main session, e.g., to exchange the business cards. When the main session is terminated, the TMs could specify whether the subordinate sessions should also be terminated (as in the Unix standard parent–child process model), or should be detached from the parent and allowed to continue. In general, the ability to manage relationships among calls is a useful feature. In the context of billing, for instance, it makes it possible to bill together calls that belong to the same high-level session, possibly at a discount.

To conclude, we note that several optimizations can be implemented in order to potentially reduce the number of messages that need to be exchanged both at call-setup time and in mid-session. For instance, the location of a downloadable service can be sent along with the initial invitation. Also, business cards can be sent as part of the mid-session signaling, as part of a SIP “info” message. However, because we try to maintain protocol-independence throughout our model, we have chosen not to implement some of these optimizations, which rely on assumptions about the underlying protocol (namely, SIP in this example).

5. Object model and implementation

In this section, we provide some details of the existing prototype implementation of CitiTime.

We describe the implementation using the UML [6], a graphical language that is rapidly gaining acceptance as a standard for modeling and designing object-oriented systems. The UML model described here was prepared using the rational rose CASE tool [13].

A UML object model as shown in Fig. 8 describes the static view of a system in terms of classes and their associations. In the figure, the boxes denote classes and the undirected lines denote associations between them. Some of the

methods available for the classes may be specified in the class box. An association can be labeled with its multiplicity at either end of the line indicating, for example, that a call may have zero or more connections (labeled as 0..*) but has a one-to-one association with a terminal connection (shown as 1..1). A specialization association, indicated by the triangle symbol, specifies that the class(es) at the bottom inherit the behavior of the single class at the top and specialize it in some way, typically by overriding some of the method definitions or by defining new methods (the new methods are usually listed in the class box.). For instance, a “DataTerminal” is a specialization of a generic “terminal”. In this figure, it should be easy to identify the classes in the JTAPI call model described earlier, and in fact each JTAPI class is represented by one class in the model.³ The purpose of the call observer class is to receive and handle events emitted by a call object. In the CCC event model, these events usually represent transitions in the state of the call. Examples of events include a new incoming call, a hang up, call answer or reject, etc. By subscribing to these events, call observers can monitor the progress of a call. When they are defined as part of the TerminalManager, they are used to inform the user on the state of the call and to ensure a consistent behavior for the TM.

CCC specialization’s to JTAPI include the PropertyManagerAddress, which is just an address with the additional capability to query a remote peer for the value of its properties (e.g., a property could be the IP address of one of its attached services, for instance, a listener object for a chat service). Similarly, the DataTerminal class is a specialization of the JTAPI terminal class that supports the notion of a simple data line (a different specialization could be for voice terminal).

In Fig. 9, we present part of a UML *sequence diagram* that shows the steps involved during a *remote call*, i.e., when a call is placed from a terminal on one provider to a terminal on a different provider, normally situated on a different machine. These steps in general subsume those involved in a

³ Some of the classes are omitted from the figure for clarity.

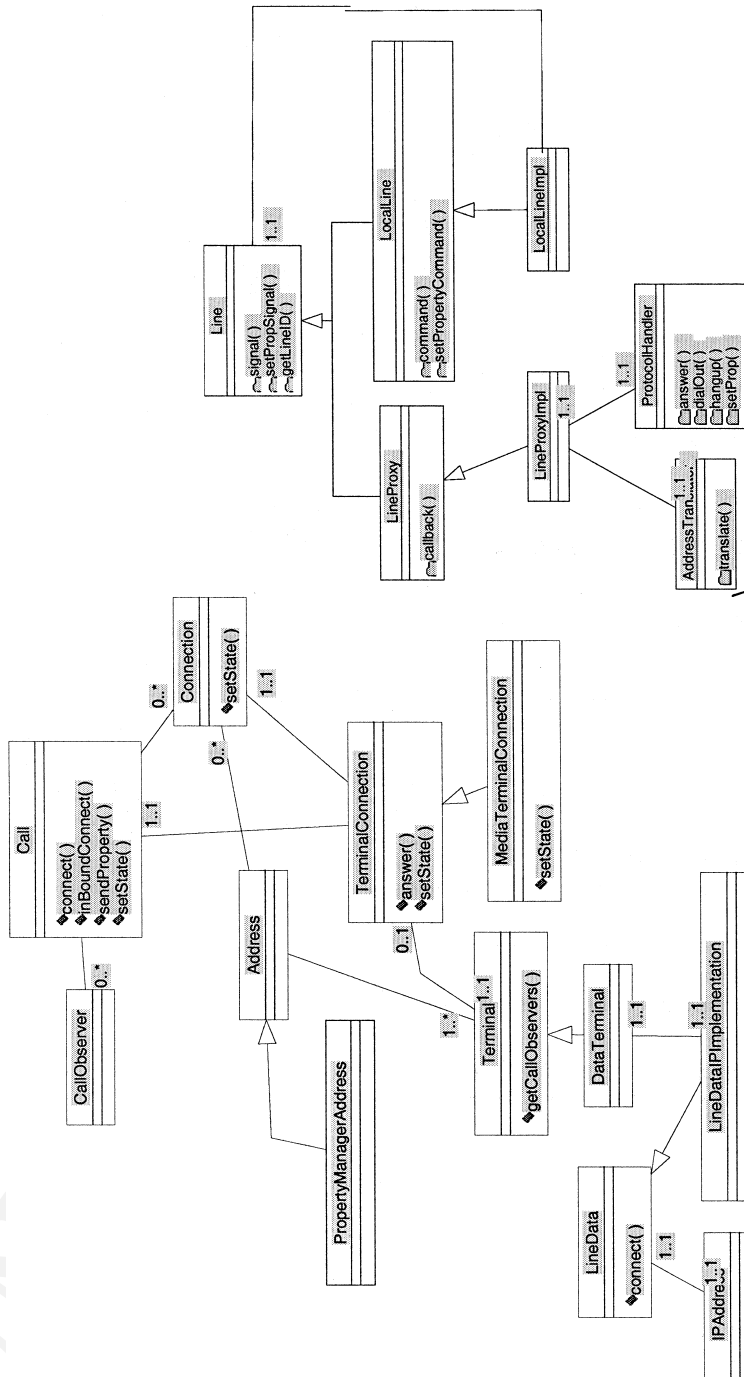


Fig. 8. Top-level and line-level object model for CitiTime.

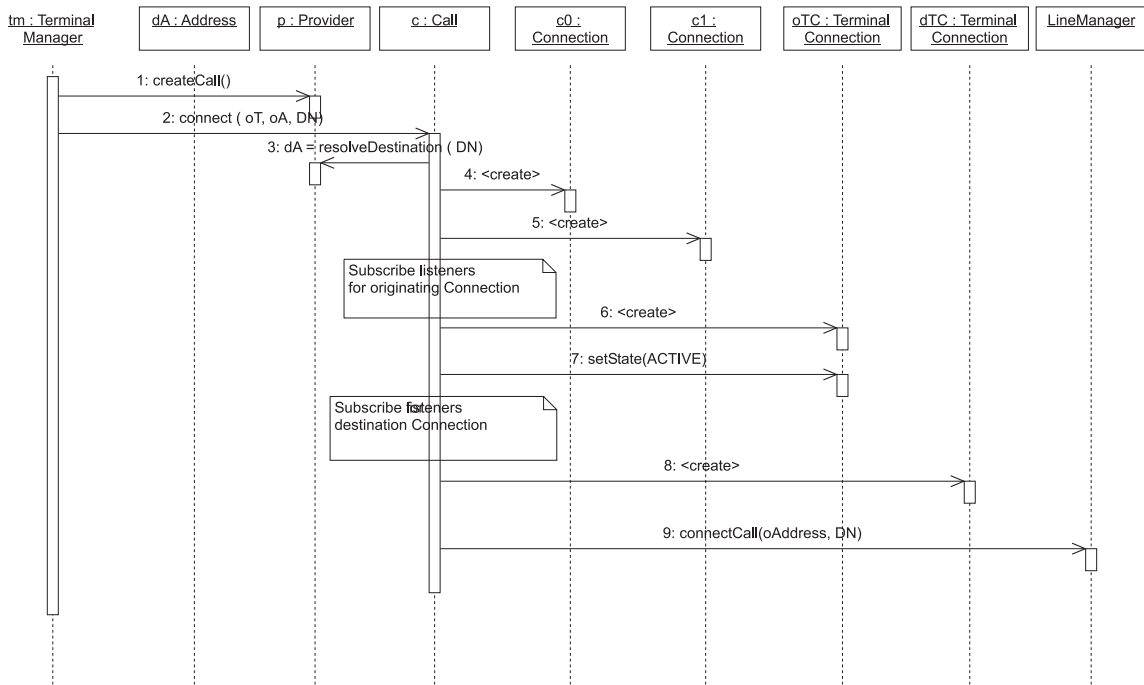


Fig. 9. Sequence diagram for remote call setup.

local call between terminals supported by the same provider. In our case, a remote call proceeds over an IP network.

Unlike the static object model, a UML sequence diagram captures the dynamics of a particular interaction among objects in a system, as it unfolds in time. The columns represent named class instances, labeled in the boxes along the top of the diagram according to their classes. The arrows represent method invocations, with their input arguments, and their return values (in many cases only the critical parameters of interest are shown). Each arrow is labeled with a number, indicating its place in the sequence. Time proceeds downwards along the page. Note that only the invocation of a method by one object on another is explicitly shown; the return is left implicit. The boxes with a dog-eared corner are simply notes explaining or providing clarifying comments. For further details on UML, see [6].

For the sequence diagram in Fig. 9, the call begins when the user enters the destination phone number and presses the “call” button on the Cit-

iTime client GUI. The TM, object residing at the client device issues a *createCall()* method invocation to the provider object, which creates a new call object and returns a reference to it. The TM then contacts the call object to set up the call, supplying the originating terminal id *oT*, the originating address *oA*, and the *DN*. The call object obtains the destination address *dA* from the provider. At this point, the *dA* is still a reference to a logical address, i.e., to a JTAPI address object. It is the task of the lower CCC layers to handle it appropriately. If the call is local, the actual *dA* object exists in the process space of the same provider, and it can be contacted directly. In the case of a remote call, the address is translated by the LineManager during the *connectCall()* method into the actual IP address and port of the destination terminal (description of the directory lookup process involved is outside the scope of the present paper.). These differences are encapsulated in the LineManager and are transparent to the high-level call object. The connection objects representing the originating and terminating legs of

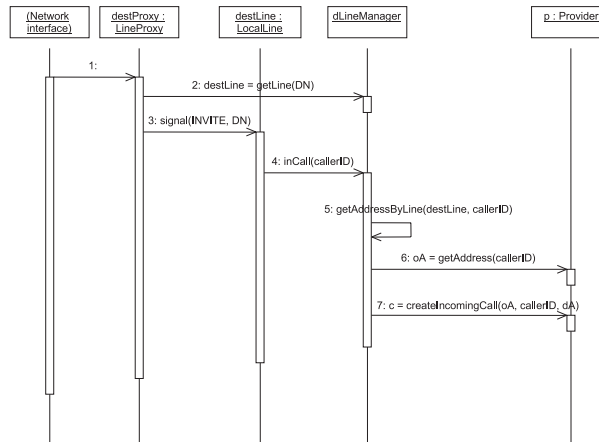


Fig. 10. Remote call-receiving side.

the call are then created, and the CallObservers defined for the call are subscribed to the call’s and connections’ events.

Note that the terminating connection *cl* in this case is simply a proxy for the connection object that will be created at the destination machine once the call setup signaling messages arrive at the destination. The use of proxies on both sides of a remote connection shields the higher layers from the details of remote signaling, by exposing a uniform behavior regardless of the type of the call. In a similar vein, when the call object creates the *oTC* and *dTC*, the latter is really a proxy for the corresponding remote object. Once the call model is setup, the call object contacts the LineManager to actually establish the connection through the signaling protocol.

The LineManager sets up a pair of line objects, representing the two legs of the call, on which the actual signaling is carried out. The handler for the signaling protocol, in this case SIP, is entirely encapsulated in the line objects. In fact, in the spirit of abstraction, from the LineManager’s perspective it is irrelevant whether the line objects are local or remote, and which signaling protocol they use. Externally, a line object for a remote destination behaves the same as for a local one. Internally, it is really implemented as a proxy, and the SIP signals to its remote peer are sent over sockets.

Fig. 10 illustrates the corresponding low-level sequence on the terminating side. A proxy for the sending line receives the initial INVITE signal from its peer through the network interface, obtains a reference to a local line that is responsible for the indicated DN, and then signals that line. At this point, the LineManager is responsible for translating the line into a JTAPI Address object that can be notified, with the help of the provider, of the incoming call.

6. Discussion and further work

In this section, we briefly discuss various aspects of the CitiTime implementation and features.

Use of SIP. We use SIP [3] as the common signaling protocol among distributed providers. Compared to other proposed protocols, such as H.323, SIP offers the advantages of specification and implementation simplicity, extensibility, and neutrality with respect both to the transport protocol and to address format. The protocol is text-based, much in the spirit of HTTP, and is designed to be part of a larger suite of session management protocols for the Internet that includes for instance the service location protocol (SLP) as well as the session definition protocol (SDP). Its simplicity allowed us to quickly implement the slight modifications we needed in order to support ne-

gotiation during call setup. Namely, we used SIP's features to define specific messages for suspending and resuming a session, for suggesting service providers' URLs, and for informing a party of the unavailability of a resource.

Deployment model. The CitiTime architecture can be deployed in various ways to provide for scalability and, potentially, for fault tolerance. The two basic deployment scenarios can be shown using the configurations depicted for IT and CTI in Fig. 3. In the distributed deployment (left side of the figure), each provider is responsible for a partition of the global address space, and providers communicate using SIP as their common signaling protocol. In the centralized, "PBX-like" configuration, one single provider manages all the connections, as well as the interface with the PSTN. Our laboratory prototype runs in both configurations at present. Clearly, other intermediate configurations are possible, ranging from applications that run on top of a dedicated provider, to one centralized call center with only one provider. Notice that providers and endpoints can themselves be connected remotely, or they can be co-located on a host.

Performance and scalability. The main focus of the CCC architecture is on its distributed control features and the flexibility in service deployment and configuration. The features do not seem to incur any penalty in terms of performance: although no specific benchmarks have been performed on the CCC architecture, simple considerations based on its distributed nature suggest that the latency experienced by the user should be comparable to that expected in the IN and CTI domains.

Thanks to the distributed nature of CCC, CitiTime can scale reasonably well with the size of the address space, i.e., ultimately, with the number of endpoints in the network. In fact, while one provider manages an entire partition over the global address space, the partitioning itself is quite arbitrary. This makes it possible to introduce new providers and to reassign addresses to providers as needed in order to balance the load across all of them. In addition, by detaching the TMs from their provider, one provider space can scale with the number of managed TMs. Address resolution

currently remains a centralized operation that can limit scalability. In CitiTime, address resolution is done simply by mapping a logical address to the IP addresses and port number of the provider responsible for that address. The mapping table is managed by a translator server that is, in principle, distinct from all of the providers. A call is setup by contacting the involved providers through the IP address returned by the translator. Each provider will then internally route the call to the appropriate TM, by setting up the call model as described earlier. Several solutions are available for scaling the translation step to a large number of addresses, such as hierarchical resolution and local caching of the translation tables. We will investigate these issues in further work.

Another clear benefit of provider distribution is its potential for fault tolerance. We can take advantage of the fact that each call object contains the full call model, i.e., it models all the legs of a call, to provide replication across providers for every single call. In the case of a single provider (i.e., of a "local call"), a detached hot-standby provider can be used for replication. CitiTime does not at present support fault tolerance, but we plan to address this issue in further work.

Finally, we mention that we have not considered the extremely important area of security in this work. We recognize this problem and solutions are being developed in the NGN space (e.g., in the Parlay framework) that may be applicable. We plan to address this in future work.

7. Related work

Recently there has been a surge of interest in programmable networks and open APIs for next-generation converged networks, i.e., networks that span the PSTN, IP and wireless networks. The JAIN and Parlay industry consortia as well as the Softswitch Consortium are currently defining standards to enable third-party service creation on software-based switches (called softswitches or call agents) or application servers attached to those switches.

8. Concluding remarks

The CitiTime system is in the process of ongoing development. A prototype implementation is currently operational in our laboratory and is being used for voice, text, whiteboard, and image-exchange sessions, as well as for advanced services such as dynamic service download.

We are continuing to experiment with the use of third-party JavaBeans software components both in the TM and the provider. Our current experience suggests that while Java and JavaBeans provide a level of portability and name standardization, an additional level of standardization is required to allow seamless plug in of software components at run time. The current CitiTime implementation overcomes this limitation by using a component launcher in addition to the component directory which manages the details for launching components. To some extent, this represents an industry-wide issue that needs to be addressed via standardization. We are also currently involved in instrumenting and configuring CitiTime for performance measurements and for evaluating its scalability.

Acknowledgements

We thank the anonymous referees for their helpful comments.

References

- [1] A. Lazar, Programming telecommunication networks, IEEE Network 11 (5) (1997) 8–18.
- [2] E. Margulies (Ed.), Understanding Java Telephony, Flatiron Publishing, 1997.
- [3] M. Handley, H. Schulzrinne, E. Schooler, SIP: Session Initiation Protocol, IETF Draft, draft-ietf-mmusic-sip-04.txt, 11 November 1997.
- [4] R. Jain, F. Anjum, P. Missier, S. Shastry, Java call control, coordination and transactions, IEEE Communications, January 2000.
- [5] JAIN Web site: <http://www.java.sun.com/jcp> and <http://www.java.sun.com/products/jain>.
- [6] C. Larman, Applying UML and Patterns, Prentice-Hall, Englewood Cliffs, NJ, 1998.
- [7] R. Englander, Understanding Java Beans, O'Reilly, 1997.
- [8] R. Orfali, D. Harkey, Client/server programming with JAVA and CORBA, second ed., Wiley, New York, 1998.
- [9] J. Nelson, CorbaBeans Proof-of-Concept, <http://www.DistributedObjects.com/>.
- [10] Sun Microsystems, Servlet Tutorial, Available from http://www.java.sun.com/products/jdk/1.2/cast-out/servlet/servlet_tutorial.html.
- [11] A. Thomas, Enterprise JavaBeans: server component model for Java, Patricia Seybold Group White Paper, December 1997, available from http://www.java.sun.com/products/ejb/white_paper.html.
- [12] Internet Mail Consortium, vCard: The Electronic Business Card Version 2.1, versit Consortium White Paper, 1 January 1997.
- [13] Rational Software Corp., Rational Rose/C++ software version 4.0.3, 1996.

Farooq Anjum is a research scientist at Telcordia Technologies, USA. His areas of research include middleware technologies, software agents, call control, mobile and wireless networking, intrusion tolerance, etc. He is also active in standardization forums like Java APIs for advanced integrated networks (JAIN). Farooq joined Telcordia after completing a Ph.D. in Electrical and Computer Engineering from the University of Maryland at College Park in 1999.



Francesco Caruso is a research scientist at Telcordia Technologies in the Information Technology and Internet Applications Department. His current interests include: distributed systems, web and Internet technologies, component architectures, and database management systems. Mr. Caruso earned a University Degree in Computer Science from University of Pisa, Italy in 1992.



Ravi Jain received the Ph.D. in Computer Science from the University of Texas at Austin in 1992. Currently he is the Director of the Middleware and Mobile Applications Research Group at Applied Research, Telcordia Technologies. His interests include programmability, middleware and applications for next-generation networks, mobile Internet access and applications, and mobile and wireless networking. Jain has numerous publications in these areas and several issued and pending patents. He is also Edit Lead for the

JAIN expert group defining Java call control (JCC). Jain is a member of the Upsilon Pi Epsilon and Phi Kappa Phi honorary societies, a senior member of IEEE, and a member of ACM.



Paolo Missier is a research scientist at Telcordia Technologies, USA. He has worked in several areas, including next-generation services for integrated telecommunication networks, large-scale and component-based software engineering, distributed object-oriented software architectures, and federated database technology. He holds an M.Sc. in Computer Science from Università di Udine, Italy and an M.Sc. in Computer Science from University of Houston, Texas, USA.



Adalberto Zordan is a consultant for Telcordia Technologies in the area of component-based distributed software architectures, and integration and migration of information systems. He holds an M.Sc. degree in Electronic Engineering from University of Padova, Italy.

UNCORRECTED PROOF