



The impact of the Internet on telecommunication architectures

Jean-Pierre Hubaux^{*}, Constant Gbaguidi¹, Shawn Koppenhoefer²,
Jean-Yves Le Boudec³

Institute for Computer Communications and Applications (ICA), Swiss Federal Institute of Technology (EPFL), Lausanne, Switzerland

Abstract

The ever-growing popularity of the Internet is dramatically changing the landscape of the communications marketplace. The two separate worlds of the Internet and Telecommunications are converging. The respective advantages of the two environments are being integrated to fulfill the promise of the information super-highways. In this paper, we examine the impact of the Internet on the main telecommunication architectures, namely the IN, the TMN and TINA. There are two new tendencies for implementing telephony services in combination with the Internet: running part of the control system over the Internet, or conveying both the user data and the control information over the Internet. We examine these two tendencies, and elaborate on possible ways of salvaging the best parts of the work achieved by the TINA-Consortium in the Internet context. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Intelligent Network; Telecommunications management network; Telecommunications information networking architecture; PSTN/Internet interworking; Internet telephony; Virtual private network

1. Introduction

Until the late 1970's, introducing new services in a telephone network was a slow and difficult process, requiring the modification of switch's software spread all over the network. Customers might not have been happy with their dependence on their operator but they had little choice in the matter. The operators may not have been happy with their dependence on their equipment suppliers, but they had no choice. At the same time, the equipment heterogene-

ity, and the absence of any consistent way to control it, caused a wealth of difficulties for anyone trying to monitor, analyze, and plan the operation of the network resources needed to provide services. This dismal situation was intolerable. Two cleansing waves simultaneously washed away some of this dependence. The first wave, called the Intelligent Network (IN) [1,2], provided a means to decouple *service-specific software* from *basic call processing*. The second wave, called Telecommunications Management Network (TMN) [3], organized, across a certain number of layers, the management of telecommunication equipment, thus facilitating the management of heterogeneous networks.

Due to the increased pace of developments in the area of Telecommunications, it was quickly perceived that, despite many of the good ideas implemented in the IN and TMN architectures alone,

^{*} Corresponding author. Tel.: +41-21-693-2627; fax: +41-21-693-6610; e-mail: jean-pierre.hubaux@epfl.ch.

¹ E-mail: constant.gbaguidi@epfl.ch.

² E-mail: shawn.koppenhoefer@epfl.ch.

³ E-mail: jean-yves.leboudec@epfl.ch.

neither one, nor even their coexistence, could address *all* the challenges of service provisioning and network management that consumers began to pose in the early 1990's. However, by merging the best ideas of those technologies, a new architecture (TINA – Telecommunications Information Networking Architecture [4]) was created by a worldwide consortium, TINA-C, formed in 1992. The Telecommunications community finally believed that they had a reference architecture for many years to come (or at least the late 1990's). Was this belief just a mirage, or a foreseen bridge to another realm?

Now, we are witnessing a third wave that is challenging the positions held by IN, TMN and even TINA: the *Internet*⁴. In the same way as the best ideas of IN and TMN led to TINA, the interweaving of the Internet with today's telecommunication solutions is generating something new that we are only beginning to discover. In recognition of some of the strengths of the Internet (e.g., open service creation environment, ubiquitous access, scalable resource location, and standardized resource access), some services traditionally based on the Public Switched Telephone Network (PSTN) are now being "ported" to the Internet. In a similar fashion, the strengths of IN, TMN, and TINA might have a positive influence on the future of the Internet. By bringing together the advantages of both worlds, it is conceivable that the disadvantages of each can be greatly diminished. This is the issue that we investigate in this paper.

We organize the paper as follows: We introduce the IN, TMN and TINA architectures (Section 2). Then, we show how the advantages of the two environments (i.e., the PSTN and the Internet) can be used to create, deploy, control and manage communication services in a more satisfactory way (Section 3). In Section 4, we examine a number of assumptions that have changed since the introduction of TINA. We show how this has changed the rules of the game in such a way that "TINA as a solution" should now be questioned; we also show that some of the better concepts from TINA could be salvaged.

⁴ In this paper, by "Internet" we designate both the worldwide network itself, and the underlying technology, namely TCP/IP (Transmission Control Protocol/Internet Protocol).

2. A reminder on telecommunication architectures

While readers familiar with IN, TMN and TINA may consider skipping the more tutorial Section 2, we describe some original concepts with Fig. 2 and Fig. 4.

2.1. The intelligent network

The term Intelligent Network (IN) was first introduced by Bellcore in the 1980's following the deployment of "800" number services in the US. The IN is an architectural concept allowing a rapid, smooth and easy introduction of new telecommunication services in the network. It was standardized by the Telecommunications Sector of the International Telecommunication Union (ITU-T), and the European Telecommunications Standards Institute (ET-SI) [5]. Detailed presentations of the IN are provided in [1,2].

The IN is based on a simple principle: the service-specific software is separated from the basic call processing and is run on a specialized node. In IN parlance, this specialized node is called the *Service Control Point* (SCP), while the switches are called *Service Switching Points* (SSP) (Fig. 1). The Service Management System (SMS) contains functions to enable the management of the IN system. The Service Data Point (SDP) is a database that holds information related to the service (e.g., customer profile). The IN is provided with a Service Creation Environment (SCE) that contains service development tools. Communication among those elements is done through the Common Channel Signalling no. 7 (CCS7) [6].

Consider the implementation of Freephone. When the SSP detects the "1-800" pattern, it temporarily suspends the basic call procedure and requests the collaboration of the SCP. The SCP realizes that the communication fees must be charged to the called party, and asks the SDP for the physical address of this party. The response from the SDP is forwarded to the SSP which can then request the connection to be set up between the two parties.

The development of the IN has been organized in several phases, each of them being characterized by

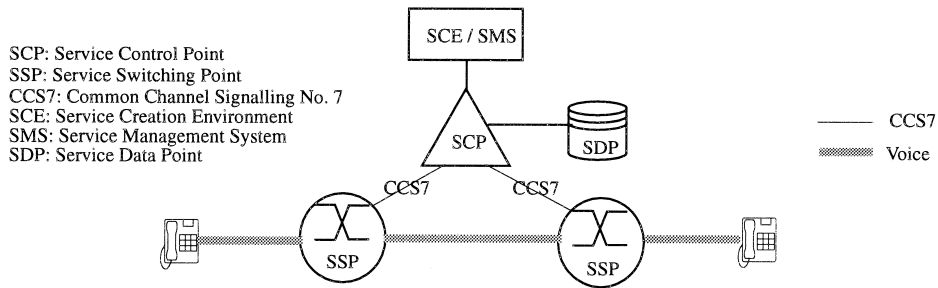


Fig. 1. Simplified architecture of the intelligent network.

a specific capability set. The oldest one is Capability Set 1 (CS-1), which covers single-ended services, i.e., service that apply to one and only one party in a call and are independent of any other parties that may be participating in the call. Examples are number translation services (e.g., call forwarding, or universal personal telecommunication), alternate billing services (e.g., credit card calling), and screening services (e.g., originating/terminating call screening, or security screening) [2].

The CS-1 is followed by two other phases (i.e., CS-2 and CS-3 [2]) which extend it with more sophisticated capabilities [2]. Much of the discussion in this article addresses only CS-1.

2.2. The telecommunications management network

The Telecommunications Management Network (TMN) [3] can be looked at from three points of view: information architecture, functional architecture, and physical architecture.

The *TMN information architecture* provides data representation of the network resources for the purpose of monitoring, control and management. The object-oriented approach is considered for the specification of the information model.

The *TMN functional architecture* describes the realization of a TMN in terms of different categories of *function blocks* and *reference points* among these blocks.

The *TMN physical architecture* corresponds to the physical realization of the functional architecture. Each function block becomes a *physical block*, and *reference points* are transformed into interfaces. The Operation System (OS) is an important physical block. It deploys the logic necessary for managing

the telecommunication activities. Among all different interfaces, the main ones are the Q3 interface (lying, e.g., between an OS and the managed resource, or between two OSs of a given management domain) and the X interface (lying between two OSs belonging to different TMN domains).

Although the TMN was defined with network management in mind, it can be used to provide a wealth of other services. One of them is the Virtual Private Network (VPN). The VPN is a telecommunication service that provides corporate networking between geographically dispersed customer premises; it is based on a shared public switched network infrastructure. The main features, which made up the early VPN specification intended for implementation over the IN infrastructure, were Private Numbering Plan (PNP), Closed User Group (CUG), and security [5,7]. PNP is a feature that enables the customer to define her own phonebook. CUG allows the customer to define groups of users; for example, there may be one CUG for top managers, another one for researchers, and so forth.

Fig. 2 shows the physical architecture of a VPN configuration management system [8,9]. The configuration management architecture consists of a set of OSs: (1) the CPN OS that manages the CPN resources, (2) the PN OS that manages the public network resources, (3) the PN-service OS that is responsible for the management of the services offered over the public network (e.g., a virtual path service in an ATM network), (4) the CPN-service OS whose role is to administer the services provided over the CPN, and finally (5) the VPN-service OS that manages the overall VPN service. The X interface enables interactions among the VPN service actors (i.e., the customer, the service provider and

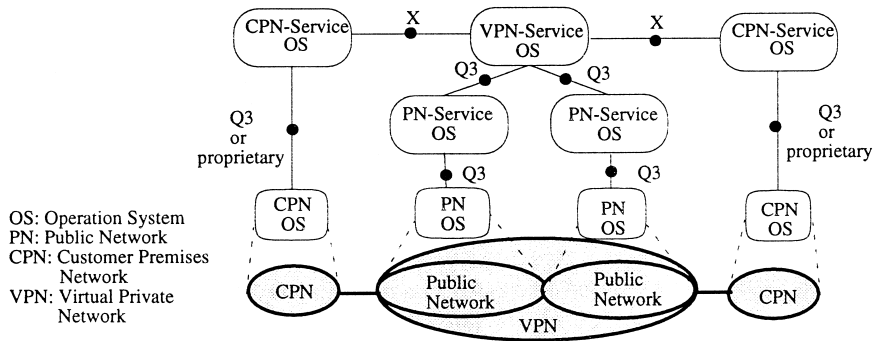


Fig. 2. VPN physical management architecture.

the network provider). The Q3 interface lies between OSs of a given management domain.

The IN and TMN architectures overlap [10]. This is illustrated by the VPN example we have just presented. Indeed, the implementation of the VPN makes use of management primitives (e.g., bandwidth reservation within the network by means of management techniques) and therefore has to be realized in the TMN. However, what happens if the VPN has to interwork with services supported by the IN, such as call forwarding or three-way calls? This example shows that unless both IN and TMN architectures are made more consistent, the interworking of IN and TMN applications will be very difficult; interoperation between applications, developed within two independent architectures, may be trou-

blesome. The TINA architecture was expected to solve this complex problem.

2.3. The telecommunications information networking architecture

A major design principle for TINA was the use of distributed computing (e.g., Open Distributed Processing, ODP [11] and the Common Object Request Broker Architecture, CORBA [12]). The rationale behind this principle was the fact that architectures based on centralized computing scale badly when the number of services and service sessions increases. A worldwide consortium was set up in 1992, named the TINA-Consortium or TINA-C [4].

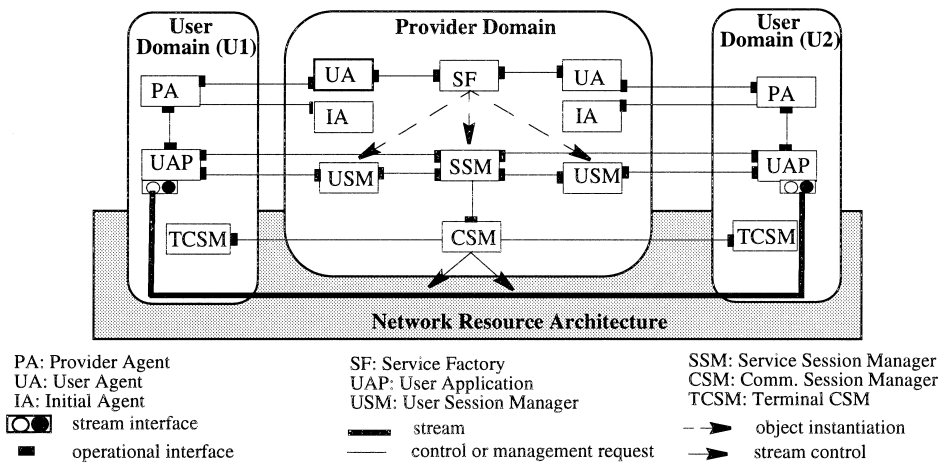


Fig. 3. TINA objects involved in a service scenario.

TINA is composed of three architectures:

- The *computing architecture* defines the concepts for designing and building distributed software and the software support environment, based on object-oriented principles. It describes the TINA Distributed Processing Environment (DPE) [13], which was recently recommended to be CORBA [14].
- The *service architecture* provides a set of concepts to build, deploy and manage TINA services [15]. In a TINA system, a service consists of a set of components interacting with one another, and deployed over the DPE. The service architecture defines the objects required for the realization of a service, their composition and interactions. The service architecture (Fig. 3) can be illustrated by considering the supervision of a multimedia communication between two end-users. In this example, an audio/video session is established between two objects corresponding to the two user applications. Other objects, located in the service provider domain, supervise the session.
- The *network resource architecture* defines a reusable group of functions which manage network transport resources and, thus, provide connection management that assists telecommunication services (distributed applications) in their need for connections.

A simplified representation of TINA is shown in Fig. 4. The Transport Network, typically based on ATM, contains the switching and transmission equipment in charge of the transportation of the user data. It is controlled by what we call in this paper the *Information Network*, which is composed of the *kernel Transport Network* (kTN) in TINA parlance, and a set of applications implemented on top of the kTN. The kTN is a network of nodes over which CORBA is deployed.

After this brief overview of the telecommunication architectures, the next two sections are devoted to the study of alternatives in the realization of these architectures using the Internet. We start with telephony services; they are basic services, from the understanding of which more complex applications can be implemented.

3. Using the Internet to unshackle telephony services

In this section, we present many advantages that can be drawn from bridging the Internet and the Telecommunications worlds. We emphasize the benefits of using the Internet for service creation. Hence, we look at the porting of the main telecommunication architecture for service creation in PSTNs (i.e., the IN) to the Internet.

Briefly, the strengths of the IN are: (1) legacy gained from the deployment of a large number of services, (2) service creation using a simple and small set of reusable Service Independent building Blocks (SIB), (3) independence of the service (and the control thereof) from the network infrastructure, and (4) a thorough conceptual model for service creation which shows a gradual “shrinking” of a service’s complexity across the model’s planes.

The current weaknesses of the IN are: (1) limited scalability: introducing a new service is very difficult because there is no standard service creation environment, the call models used by the switches in the network are not the same, and service interaction [16,17] is still a challenging problem, and (2) limited openness: interoperability between IN systems operated by different stakeholders is still an issue to address and there are significant differences between IN architectures (e.g., the Advanced Intelligent Net-




Fig. 4. A simplified view of TINA.

work from Bellcore uses a protocol and a call model that are incompatible with the ones recommended by the ETSI).

Regarding the Internet, its main strengths are [18]: (1) an open service creation environment for content services, (2) a devolved service architecture management: each user can create and manage her own service profile, for example a Web page, (3) ubiquitous access via the Internet protocol stack implemented world-wide, (4) hierarchical, scalable translation service through a powerful Domain Name System (DNS), (5) scalable resource location (supporting over 20 million URLs today), (6) standardized resource access (HTTP, CGI, MIME), (7) standardized service logic languages (HTML, Java), and (8) very fast service deployment environment (Web pages and CGI scripts easily written).

The current weaknesses of the Internet are: (1) relatively low level of security (even if this level is increasing, the public is still skeptical), (2) only best-effort quality of service in the current implementation, (3) little commitment of Internet Service Providers (ISP) to guarantee quality of service within and across their domain.

There are *two tendencies* toward bringing together the advantages of the Internet and the Telecommunications worlds. The *first* tendency uses the PSTN as the medium transport infrastructure, but devolves part of the request mechanisms to be run on the Internet. This tendency is exemplified by the work going on within the PSTN/Internet Interworking (PINT) group [19] at the Internet Engineering Task Force (IETF). The PINT contribution is described in Section 3.1.

The *second* tendency alleviates the PSTN as the principal infrastructure for the transport of user data, and promotes the Internet as a serious alternative for conveying this data. As a result, the complete telephony service architecture will be deployed over the Internet: this is IP telephony (Section 3.2).

Both tendencies are presented in the following sections, and some of their strengths and weaknesses are revealed in the context of the Virtual Private Network (VPN) which is historically one of the services that first demonstrated the limitations of the IN model in providing fairly complex services. These limitations are mainly due to the lack of interworking between IN systems operated by different

providers, and the lack of standards related to the management of the IN [20].

3.1. Interworking between the PSTN and the Internet

In this section, we present the general framework laid out by the IETF PINT working group (Section 3.1.1). Then, we report on realizations that fit into the PINT reference model, particularly the WebIN architecture [18,21] (Section 3.1.2). The focus of the PINT working group is on enabling the user to *request* IN-supported services from the Internet. We shall see that WebIN deploys a higher number of *control* mechanisms (not only request mechanisms) on the Internet than intended by the PINT group. Lastly, Section 3.1.3 presents a possible implementation of the VPN over a combination of PSTN and Internet.

3.1.1. The PINT reference model

The IETF PINT group suggested a way that services currently provided over PSTNs could take advantage of the strengths of the Internet. They proposed a framework which achieves this in [19] (updated in [22]). This on-going work emphasizes the interfaces between the modules that make up the framework (Fig. 5). The main modules are: Service Nodes (SN), Service Management Systems (SMS), Web Servers, PSTN/Internet Gateways, Service Control Points (SCP), Central Offices, and Mobile Switching Centers (MSC). The main difference between an SN and an SCP is that, unlike the latter, SNs can perform switching functions as well as activities traditionally carried out by specialized resources (such as voice recognition devices, and text to audio transcoders). In simpler terms, if switching or specialized functions are needed, then SNs must be used instead of SCPs.

Four services are considered as case studies for the reference model in Fig. 5. These services are: Click-to-Dial, Click-to-Fax, Click-to-Fax-Back, and Voice-Access-to-Content. Consider the usage and implementation of the Click-to-Dial service. While browsing the Web page of a company, a user becomes interested in talking to a sales representative. To do so, she simply clicks on a button to send her request. The user is understood to request the phone connection to be set up between her telephone equip-

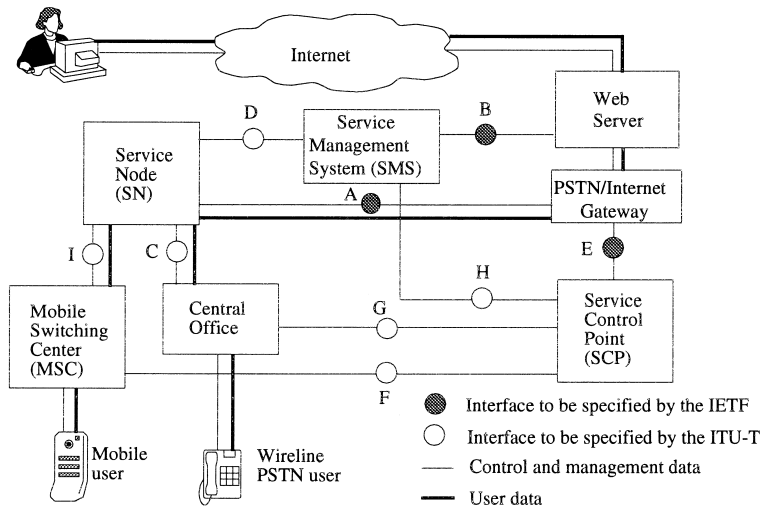


Fig. 5. Reference model for PSTN/Internet Interworking (PINT) (augmented from [19]).

ment and the one of the sales representative. The Web server of the company concerned forwards the request to a PSTN/Internet gateway which, in turn, calls the Service Node (SN) through an A interface (Fig. 5). Upon receipt of the request, the SN determines which sales representative to call; the company may have subscribed to a call distribution (or routing) service. Afterwards, the SN passes the chosen representative's address to the central office—or the mobile switching center—to initiate the call. This is done through either the C or I interface. The company can update its profile by interacting, via the Web server, with the SMS through a B interface which is based on the Simple Network Management Protocol (SNMP) [23].

The thorough description of the interfaces between the components of the PINT reference model is under specification. These interfaces are cast with respect to their relevance to the standards bodies at play, namely the IETF and the ITU. The interfaces A, B, and E would be standardized by the IETF, while the ITU would deal with the rest [24].

Many implementation efforts were carried out in the area of Internet/PSTN interworking. In the following section, we report on some of them.

3.1.2. Realization of Internet / PSTN interworking

A great deal of effort was expended in the past on the design of gateways between the PSTN and the

Internet. This effort was spent by several companies like Siemens [25], Hewlett-Packard [21], IBM [26], and Vocaltec Communications [27]. None of their solutions, except that from Hewlett-Packard, addressed the particular issue of porting the IN architecture to the Internet. Consequently, we concentrate on the project from HP, which was developed at the HP laboratories in Bristol, UK. The main outcome of this project is a prototype architecture called WebIN (Fig. 6). WebIN resembles the IN standard architecture, except in three respects: (1) SDPs and the SMS are connected to the Internet; (2) SCPs are connected to SDPs through a DPE that could be either the Web (HTTP/CGI) or CORBA; (3) the role of the SCP is confined to finding out the correct reference of the SDP that corresponds to the service actually invoked.

Note that in Fig. 6 the SCE (Service Creation Environment) is located within the Internet domain. Hence, WebIN can benefit from the easy service creation environment (HTML, Java, etc.) on top of the Internet. After writing the service logic (using HTML, Java, Perl, etc.), the service creator simply uploads her service logic into the Web server dedicated to her.

3.1.3. Provision of the Virtual Private Network (VPN) service

In this section, we show how well a fairly complex service such as the VPN is currently imple-

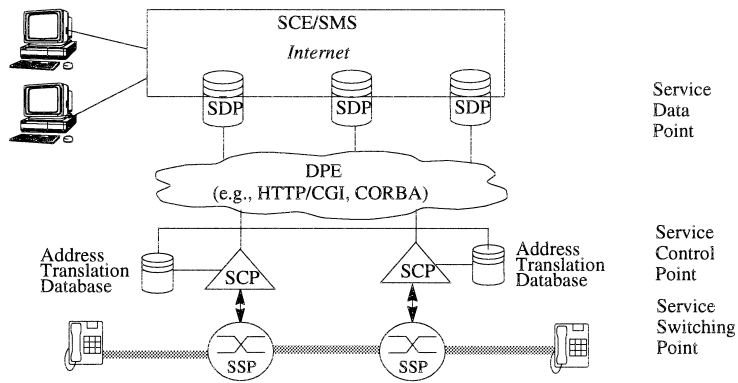


Fig. 6. The WebIN physical architecture (adapted from [28]).

mented over the PSTN, and how it could be provided using the Internet for control and management. We consider two sites located in two different countries serviced by two public operators.

3.1.3.1. VPN implementation over the PSTN. Consider user A in Lausanne, Switzerland, wishing to communicate with user B in New York, USA (Fig. 7); both users belong to the same VPN. Since two IN infrastructures from different operators rarely interoperate, the PNP and CUG descriptions must be implemented in both networks. User A uses the extension number to communicate with user B. A translation from extension numbers to physical telephone addresses is needed at the SCP connected to the SSP, which user A is attached to. The physical number of user B is then found. Furthermore, the SDP, that holds the VPN database of the company of interest, checks whether the two users belong to the same CUG and are not subject to any call restrictions

(e.g., not being able to generate calls to outside of the company, not being able to interact with people outside the CUG, etc.). After everything is checked, the SSP related to user A is instructed to generate a call request to user B's physical address. The same processing phases take place when user B wishes to communicate with user A. The IN system in the USA is understood to comply with Bellcore AIN, while the one in Europe complies with ETSI recommendations. As mentioned in the introduction of Section 3, the two IN systems considered in the example are incompatible. Therefore, the VPN profile must be duplicated in each IN infrastructure. The consistency of this profile everywhere then becomes an issue in the very likely situation where all the databases are not updated at the same time. Consider that we move user A from one CUG to another in the same VPN. If the change is not simultaneously reflected everywhere, user A will be using the privileges of the two CUGs for some time.

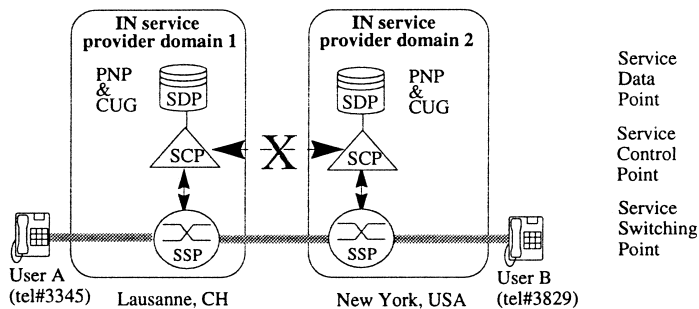


Fig. 7. Provision of a VPN service across two IN provider domains.

3.1.3.2. *Running the VPN control system on the Internet.* Fig. 8 depicts a possible realization of the VPN service in a context that uses the Internet to run the control system. The SSP associated with the caller, issues a request to the SCP with the caller’s—as well as the callee’s—identity and telephone number. Each IN operator is given the address of a Web server where the service data and logic can be requested. The Web servers attached to the SCPs cooperate through the Internet in order to find the right location of the data needed. Specifically, the data related to the CUG to which users A and B belong, may not be held by the server attached to the SCP in Lausanne. Assuming that the CUG data resides on the server attached to the SCP in New York, there will be interactions, over the Internet, between the two servers involved. Therefore, the PNP and CUG profiles can be split over the servers used, and there is no need for data duplication (unlike in the case described in Section 3.1.3.1). Changes in the VPN profile are immediately reflected. Therefore, the lack of *interoperability* between IN infrastructures is solved in a simple way in the case where the control and management system is run on the Internet. Moreover, the customer can update the VPN profile through the Internet; she therefore is equipped with *control and management capabilities* over her service. She can perform these tasks without interacting with the VPN service providers.

However, running the control system on the Internet presents a number of issues such as security and efficiency (the processing of a service request can take a longer time than in the legacy IN). Security

can be assured by using tunneling and data encryption techniques. Improving the efficiency of the scheme needs more research.

The PINT reference model described above still uses the PSTN as the medium carrier (i.e., communications still go through the PSTN). A radically different approach is taken by the promoters of telephony over the Internet. In their scheme, not only part of the control system runs on the Internet, but all the control as well as the user information transfer will be achieved using the Internet. This approach is described in the next section.

3.2. *IP telephony: the Internet as a transport network*

The tremendous market penetration of the Internet makes it possible to envisage its use as an integral voice transport network. Schulzrinne [29] assessed the situation of both the Internet and the legacy telephone system. He provided valuable motivations for re-engineering the latter. However, there are issues that undermine the provision of telephony services on the Internet: the cost (do we really want to buy a \$2000 phone, even though we may use it for other purposes?), and the delay suffered by packets traveling on the Internet. Both issues, especially the last one, are being worked on intensely, as we report below.

In Section 3.2.1, we present the overall framework for IP telephony. In Section 3.2.2, we propose a technical implementation framework that identifies the main protocols to be used when providing IP telephony. In Section 3.2.3, the provision of the VPN

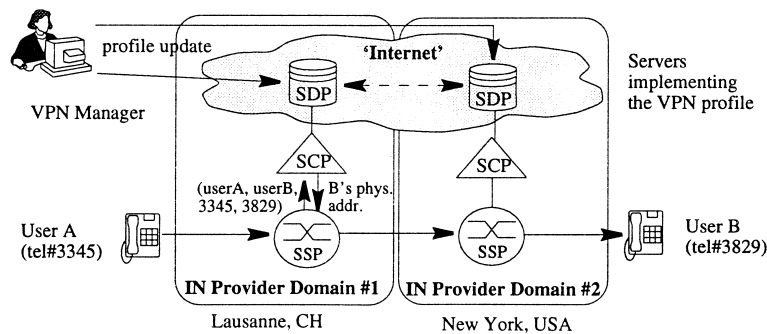


Fig. 8. A realization of the VPN service with the control system.

service within the IP telephony framework is briefly discussed.

3.2.1. Architecture

A standard framework for IP telephony (Fig. 9) was proposed by Sinnreich and Young in [30]. The main ideas behind this architecture are *continuous data flow* through the Internet (unlike the work described in Section 3.1), and *interoperability* among legacy telephone equipment and computers. Four groups of possible interactions were identified: (1) gateway-to-gateway signaling and information transport, (2) interactions between the SCP and the gateway databases, (3) PC-to-network signaling, and (4) management information flow.

3.2.2. A Technical implementation framework

Two relevant proposals for IP telephony are the ITU-T recommendation H.323 [31–33] and the Session Initiation Protocol (SIP) [34] being worked on at the IETF.

H.323's primary goal is to enable audiovisual interactions (i.e., provision and control of audio/video interactive services) among terminals situated in various environments with/without Quality of Service (QoS) guarantees. The key H.323 components that are important to our discussion here are the *gateway* and the *gatekeeper*. An H.323 gateway is in charge of translating call signaling, control messages and multiplexing techniques across the network technologies involved in the call. While the gateway is an optional element in an H.323 system,

its absence, however, undermines interoperability among differing network technologies. Therefore, H.323 fully fits within the framework depicted in Fig. 9. The internal architecture of a gateway is not prescribed in H.323; a possible layout can be found in [25]. An H.323 gatekeeper performs network administration, bandwidth re-negotiation and address translation. H.323 terminals must get permission from the gatekeeper before they can place or accept a call. A gatekeeper is not required in an H.323 system; however, its presence empowers the network administrator with much control over the network, and provides users with the possibility to define alias addresses that are independent of the network addresses. The gatekeeper is a helpful component that can be used to implement many supplementary services such as some of those currently supported by the IN (e.g., call forwarding, call screening).

SIP is a protocol that enables the invitation of users to participate in multimedia sessions. It also enables user mobility by forwarding invitations to the user's current location. A more precise description of the use of SIP for IP telephony can be found in [35]. Unlike H.323 which describes an entire system encompassing many network technologies, SIP has a more limited scope; it is mainly intended for the Internet.

There is one issue that neither H.323 nor SIP have currently addressed: the resource reservation style to be used within the Internet. This style may be any of that by the IETF Integrated Services (int-serv) group, the one proposed by the Differentiated Services




Fig. 9. A standards framework for IP telephony [30].

(diff-serv) group, the YEt another Sender Session Internet Reservation (YESSIR) protocol [36], or the Scalable Resource Reservation Protocol (SRP) [37].

The int-serv group proposes an integrated services architecture [38] that recommends the Resource reSerVation Protocol (RSVP) [39,40] as the protocol to be used for reserving resources in the Internet. Two network services can currently be requested with RSVP: guaranteed service [41] and controlled-load service [42]. The former service guarantees the traffic (burst, average rate) and reservation parameters (bandwidth, delay, and losses) negotiated at the connection setup time. Guaranteed service is required by applications that have tight quality requirements. The controlled-load service guarantees only the traffic characteristics negotiated. It can be requested by applications that can adapt to the network state.

Over the years, RSVP showed a number of drawbacks, importantly the necessity to keep the state of each and every flow that travels through the router. The state information to be kept will then grow significantly as the number of users increases. This weakness has steered new alternatives, such as

- YESSIR: an extension of the Real-time Transport Protocol (RTP) [43] to support resource reservation along the path between the sender and the receiver;
- Differentiated services [44]: a scalable resource reservation style that aggregates packets from different sources into a number of service classes. These packets are then treated according to the class to which they belong. Compared to the int-serv solution, diff-serv proposals reason in terms of service classes, instead of flows with specific source and destination addresses. In this way, the state information to be kept by the routers is proportional to the number of service classes.
- SRP is a reservation protocol which achieves scaling for a large number of flows by aggregation on each link in the network. Users do not explicitly signal connection parameters; instead, senders mark packets not covered by an existing reservation as REQUESTs. A router's decision to accept a REQUEST packet is based on an estimate from measurements of already reserved traffic.

The main strength of these alternatives over RSVP is that they demand little modification in the routers as currently implemented. Guérin [45] gives a more detailed discussion on new trends for delivery in packet networks. Moreover, an interesting critique of the Internet paradigm for real-time services is given in [46].

3.2.3. Provision of the VPN service

The provision of the VPN service within the IP telephony realm is straightforward from Fig. 8: SSPs can be replaced by QoS-capable routers, and the WebSCP can be replaced by a SIP server. The control of the VPN service will be achieved through interactions among these servers: the address of the Web site that hosts the VPN profile will be given by a SIP server.

3.3. Conclusion on telephony services

In this section, we have presented the influence of the Internet on the IN through the two main tendencies identified from the literature. These tendencies are: use of the Internet to implement part of the service control system, and consideration of the Internet as a competing medium transport infrastructure to the legacy PSTN. We assessed how well the tendencies support a fairly complex service, i.e., the VPN. We could show that interoperation among IN infrastructures can be achieved when the Internet is used to run the service control system (Fig. 8). Moreover, the customer can update the profile of his VPN without interacting with the service provider. Services can therefore be easily created.

Even though the two tendencies outlined already have major implementations, there are many issues yet to be solved. For the tendency exemplified by the PINT working group, clear interface specifications between the framework elements are still to come, whereas several inconsistency flaws infect the IP telephony framework [46]. Nevertheless, there is serious likelihood that the two solutions will be deployed on a large scale in the future.

After discussing the impact of the Internet on telephony services, we concentrate, in the next section, on the future of the TINA architecture within a context more and more in favor of the Internet.

4. An outlook for TINA

4.1. Introduction and assessment of the current situation

From its very beginning, TINA has been a project funded and carried out by the traditional telecommunication companies, namely the network operators and their equipment providers. These powerful stakeholders could (and still can) have leverage on the following factors:

- high dependability of the services they provide (wireline and wireless telephony, ISDN, data transfer over X.25, Frame Relay and Switched Multimegabit Service – SMDS);
- strong internal research effort (but sometimes with a limited awareness of what is happening “outside”);
- in most cases, strong dependence on the political power of their home country.

These factors proved to be extremely efficient even in the recent past, allowing very ambitious projects to be defined and implemented on a wide scale; one of the most convincing examples to illustrate this capacity is the successful deployment of the GSM (Global System for Mobile communications) in Europe and in many other regions around the world.

However, these factors can also become weaknesses if the rules of the game are suddenly modified, and especially if the evolution of the telecommunications business becomes dramatically unpredictable. To some extent, TINA has been a victim of this change. To better understand what happened, it is worth having a close look at some of the, explicit or implicit, assumptions on which the TINA initiative is based. It is important that previously ingrained doctrines and solutions based on it, be re-evaluated in order to determine whether they still make sense. This will help us better identify the future perspectives.

Assumption 1. The services will continue to be provided by the network or network-based servers rather than the end-user terminals.

This emphasis on the role of the network might be defended given the following arguments:

- The IN CS-1 was already partially deployed at the time TINA was launched; the IN had proved its ability to support a significant number of (voice-based) services.
- The World-Wide Web was still in its prototyping phase, and very few telecommunication companies had realized at that time its latent power.
- Network operators had an understandable reluctance to see their role reduced to that of humble connectivity providers.

In the meantime, the Internet showed certain dynamism in service provisioning, which is not quite the case of TINA: New software (e.g., Java applets [47]) can be downloaded by terminals during an ongoing session. TINA should then be revised in order to take into account the possibility for some objects (Fig. 3) in the service provider domain to be executed in the customer domain, because of the increasing processing power and “intelligence” of the end-systems.

Network operators and service providers should realize that the interface between the terminal and the network does not necessarily correspond to the management responsibility border between the service provider and the user. For instance, a service consisting in hosting Web pages (of a company) can be provided by an operator or a third party. In this case, the company outsources information and network-related services to another service actor; we believe that this is a key new way to look at service provision. Beyond outsourcing, operators should also concentrate on the provision of services that are better supported by the organization operating the connectivity means. This may be because economies of scale can be found, or because the placement of the service inside the network bears some fundamental advantages. Examples are mobility mechanisms, which allow the users to access the same services from anywhere, as well as routing mechanisms. By concentrating on these aspects, network operators can protect themselves from becoming mere connectivity providers and revitalize their role in the provision of services to end-users.

Assumption 2. The B-ISDN will be based on end-to-end ATM.

ATM was defined with the precise goal of being the transfer mode to support B-ISDN. In spite of the

fact that TINA tried to remain technology independent, the transport network of TINA was usually considered to be a high-speed and connection-oriented network. The possibility of coping with a connectionless transport network has only been taken into account recently [48].

New technologies, such as the Gigabit Ethernet [49] and Terabit Switches/Routers [50], can deliver satisfying QoS with less overall complexity than ATM. Therefore, it is likely that the B-ISDN will leverage on these technologies, at least in access networks, while ATM will mainly stand as a backbone network technology.

Assumption 3. The pace of the telecommunications evolution will remain under the control of the main telecommunication stakeholders.

This assumption is clearly based on the monopolistic (or oligopolistic) tradition of this domain; for decades, and until only recently, it has been almost impossible for a newcomer to capture a share of the Telecommunications market, at least as a network operator. Assuming that the pace of evolution can be controlled as before, reflects a nostalgia of the golden age, during which the former CCITT (now ITU-T) would issue a set of new recommendations at the end of four-year-long *study periods*. The TINA Consortium followed the same step.

Meanwhile, the development of the Internet protocols has proved that it is possible to “do first and standardize later”. *Incremental development* is now a usual practice, which unbelievably speeds up the development process.

TINA is not the only victim of the change introduced by the Internet; the legacy telephone network was also shaken out. Isenberg [51] assessed this impact by analyzing, as we did for TINA, the assumptions made by operators of the telephone network.

4.2. What can be salvaged from TINA

In the previous section, we have seen that most of the assumptions which led to TINA are no longer valid because of the complete change in landscape brought on by the Internet. However, the Internet has

not (yet) fully responded to the challenge of the provision of multiparty, multimedia services. In particular, the guarantee of a given quality of service is still problematic [46]. This means that even if TINA will never be implemented to the extent it was initially intended, some of the better concepts could be re-used. Recently, several TINA proponents have been studying the applicability of TINA to the Internet [52–55].

What can be salvaged from TINA? Certainly, the idea of deploying an Object Request Broker (ORB), such as CORBA, in order to support sophisticated and distributed supervision functions, has a strong potential. Unlike TINA-C, we believe that it is unrealistic to entrust the establishment and release of communications to the kTN, because its ORB-based philosophy cannot currently cope with the real-time nature of the problem in an efficient way. Even though progress is being made in ORBs to support real-time services, the performance achieved is still low, compared to that of typical solutions using conventional signaling or routing of IP frames, without any middleware in between. Therefore, communication establishment and release, as well as data-stream transport, are better performed by the Transport Network. Hence, it is most probably in the area of management, and management-based services, that the kTN of TINA can be useful.

A possible integration of the Internet with the TINA architecture, based on the above analysis, is shown in Fig. 10 (which is derived from Fig. 4 by replacing ATM with IP). Unlike with TINA, the establishment and release of short-lived flows, or connections, in the architecture that we propose here would be fully achieved within the Transport Network itself. This would alleviate the lack of performance observed when the connections are entrusted to the Information Network (due to the high latency of CORBA). The Information Network would be devoted to functions that require a more sophisticated (albeit slower) software machinery. This could include service management and billing, and could be based on a distributed object-oriented framework such as CORBA. This evolution is compliant with the trends in network management research [56]. Communications among the nodes of the Information Network and the routers of the Transport Network would be supported by SNMP.

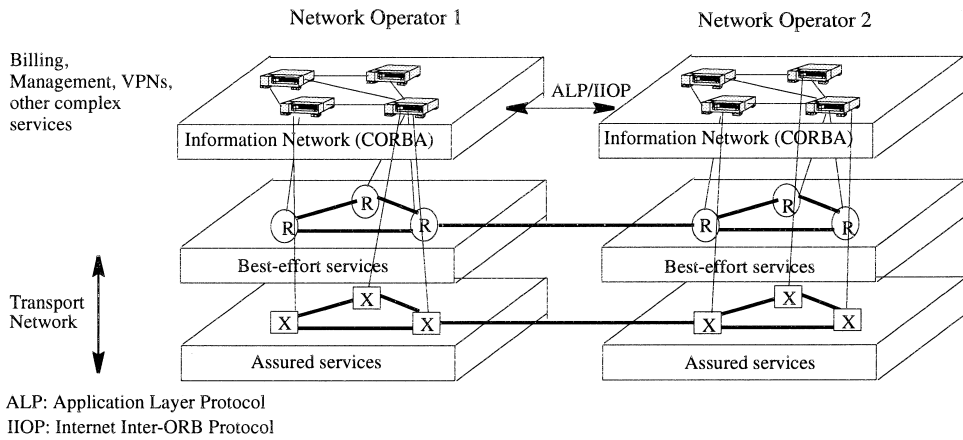


Fig. 10. A possible integration of the Internet with the TINA architecture.

Control and management functions (e.g., billing, resource management, and fault management), located within the application layer of the Information Network, interact by means of an Application Layer Protocol (ALP), which is the Common Management Information Protocol (CMIP) in the case of the OSI management model. The interface between two co-operating network provider domains is then a TMN X interface (see Section 2.2). The application pieces use the kTN to communicate with one another. The Internet Inter-ORB protocol (IIOF) supports the communication between kTNs of different network operators, and can help implement the functionality of a TMN X interface.

As for the Transport Network, it is understood to provide two categories of services: assured services and best-effort services. *We define an assured service as a service that requires some resource reservation mechanism from the network*; this service may be firm (e.g., guaranteed service considered in [41]) or looser (e.g., the controlled-load service [42]). In this case, the network operator is expected to establish connections with the requested QoS. In the case where the technology supporting the IP network is an ATM backbone, the connectivity is provided by means of Virtual Path Connections or Virtual Channel Connections. In a pure IP Wide Area Network (not based on ATM), reservations can be made by RSVP or by newer reservation styles such as differentiated services [44]. One way to optimize resource

allocation between assured and best-effort services is described in [57].

The provision of the VPN service, as defined in Section 2.2 and discussed throughout Section 3, would be easy to implement within the proposed architecture: all Operation Systems would reside in the Information Network; they would monitor the resources devoted to the assured services, in the Transport Network.

Another aspect that can be salvaged from TINA is the idea of *generic* session control. In the IETF, there is a tendency to keep creating new protocols, leading to a growing complexity of the software on the terminals; the genericity endeavor of TINA can be useful for the implementation of the numerous functions that we placed in the Information Network.

The TINA initiative can be salvaged provided:

- the 3 assumptions discussed in Section 4.1 are reconsidered,
- the importance of the Internet is recognized,
- the focus is set on the most important technical challenge: the provision of information services based on the Internet,
- the attempt to define an abstract framework from scratch is avoided, even for management services.

The second phase of the TINA project (1998–2000) [58] integrates part of our critique. The importance of the Internet is being partially taken into account, and more emphasis is being laid on the implementation of the solutions proposed.

5. Conclusion

The tight boundaries that grew up over time between the telecommunications and the Internet worlds have been falling apart for the last few years; these worlds can no longer ignore each other. The successful ideas of both communities are being brought together in order to create a new paradigm that satisfies all of the actors involved in the provision of services, using technologies from the two realms. The dependability and reliability of guaranteed services in the telecommunications world can inspire the Internet community. Conversely, the ease of service creation, deployment, control and management, that characterizes the Internet can help alleviate some of the weaknesses of the telecommunications world.

In this paper, we investigated the impact of the Internet on the main telecommunication architectures (i.e., the IN, TMN and TINA). We used the VPN as a case study to assess the influence of the Internet over these architectures.

Two tendencies were outlined for the interworking between the PSTN, over which the IN is deployed, and the Internet. The first tendency keeps the PSTN as the only transport infrastructure, but devolves part of the request mechanisms (e.g., service data and logic) to be run on the Internet. A reference model for this tendency is proposed by the IETF PINT group. We sketched an implementation of the VPN service which uses the Internet to run the control system. We then showed that interoperation between any two IN systems operated by different stakeholders can be easily achieved. This interoperation is still a problem in existing IN systems, even though the IN standards were released a long time ago. The second tendency towards using the Internet to deploy telephony services, promotes the Internet as an integral transport infrastructure: this is IP telephony.

We did not directly address the “porting” of the TMN to the Internet, since TMN is rather a framework than a real architecture with well specified elements. However, the problems that this might cause are revealed in our study of the influence of the Internet on TINA. We started with the examination of some relevant assumptions that led to the complexity of the solutions proposed by TINA. These

assumptions reflected the situation years ago, when the Internet was just that “toy network” without any potential to compete with telecommunication networks. The advent of the Web and advanced technologies such as Java brutally changed the course of things, and drove the telecommunications world into an embarrassing position. We elaborated on the TINA ideas that are likely to survive; these are the concept of generic session control and the use of ORBs. This led us to an architecture made up of two main networks: an IP-based Transport Network and a CORBA-based Information Network. The establishment of short-lived connections is entrusted to the Transport Network, while long-lived connections are enforced by the Information Network, which also supports service control and management mechanisms. The Transport Network can cope with both best-effort services and assured services.

A crucial issue is the study of the convergence of the Internet and the Telecommunications worlds. In particular, the comparison of mobility services in the two worlds, and the way they can be integrated, are key notions that we are currently investigating.

Acknowledgements

The authors are indebted to the Guest Editors and the anonymous reviewers for their relevant feedback which helped improve the quality of the article. They also wish to express their gratitude to Joergen Dyst, Tim Eckardt, Pierre-Alain Etique, Igor Faynberg, Henning Schulzrinne and Chris Smith, as well as Holly Cogliati, Sebastian Staamann, Paul Hurley, Falk Dietrich and Xavier Logean for their valuable inputs to, and comments on this paper.

References

- [1] I. Faynberg, L.R. Gabuzda, M.P. Kaplan, N.J. Shah, *The Intelligent Network Standards, their Application to Services*, McGraw-Hill, New York, 1996.
- [2] T. Magedanz, R. Popescu-Zeletin, *Intelligent Networks: Basic Technology, Standards and Evolution*, Intl. Thomson Computer Press Ed., 1996.
- [3] ITU-T, *Principles for a Telecommunications Management Network*, Rec. M.3010, May 1996.
- [4] W.J. Barr, T. Boyd, Y. Inoue, *The TINA initiative*, *IEEE Communications Magazine* (March 1993) 70–76.

- [5] ITU-T, Introduction to Intelligent Network Capability Set 1, Rec. Q.1211, March 1993.
- [6] A.R. Modaresi, R.A. Skoog, Signaling system no. 7: a tutorial, *IEEE Communications Magazine* (July 1990) 19–35.
- [7] J.-P. Gaspoz, Object oriented method and architecture for virtual private network service management, Ph.D. Dissertation No. 1446, Communication Systems Division, Swiss Federal Institute of Technology, Lausanne, Switzerland, 1995.
- [8] J.-P. Gaspoz, J.-P. Hubaux, S. Znaty, A generic architecture for VPN configuration management, *Journal of Network and Systems Management* 4 (4) (1996) 375–395.
- [9] J.-P. Gaspoz, Methodology for the development of distributed telecommunication services, *Journal of Systems and Software*, 1996.
- [10] G. Pavlou, D. Griffin, An evolutionary approach towards the future of IN and TMN, *Interoperable Communication Networks, Special Issue on Telecommunication Services*, 1 (1998) 1–16.
- [11] ITU-T, Basic Reference Model of Open Distributed Processing – Part 1: Overview and Guide to Use, Rec. X901, 1994.
- [12] OMG, The Common Object Request Broker: Architecture and Specification, Rev. 2.0, July 1995.
- [13] TINA-C, TINA Distributed Processing Environment (TINA-DPE), TINA-C Deliverable No. TB_PL001_1.0_95, July 1995.
- [14] S. Vinoski, CORBA: integrating diverse applications within distributed heterogeneous environments, *IEEE Communications Magazine* (February 1997).
- [15] TINA-C, Service Architecture, Vers. 5.0, 16 June 1997.
- [16] P.-A. Etique, J.-P. Hubaux, X. Logean, Service Specification and Validation for the Intelligent Network, *Interoperable Communication Networks, Special Issue on Telecommunication Services*, 1 (1998) 41–69.
- [17] P.-A. Etique, Service specification, verification and validation for the Intelligent Network, Ph.D. Dissertation No. 1442, Communication Systems Division, Swiss Federal Institute of Technology, Lausanne, Switzerland, 1995.
- [18] J. O’Connell, P. Hoft, Web-IN: integrating intelligent networks and the Internet, HP white paper, 30 November 1996.
- [19] I. Faynberg, M. Krishnaswamy, H. Lu, A Proposal for Internet and Public Switched Telephone Networks (PSTN) Internetworking, Internet Draft, March 1997.
- [20] RACE Project PREPARE, IBC VPN Services, Specification CFS D721, December 1993.
- [21] C. Low, Integrating communication services, *IEEE Communications Magazine* 35 (6) (1997) .
- [22] M. Krishnaswamy, PSTN-Internet Interworking – An Architecture Overview, Internet Draft, November 1997.
- [23] J. Case et al., Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2), IETF RFC 1905, Draft Standard, January 1996.
- [24] L. Conroy, M. Lautenbacher, R. Scholl, Analysis of Services and Interfaces used when Interworking the Internet and the Intelligent Network (I.N.), Internet Draft, July 1997, Work in progress.
- [25] P. Simeonov et al., @INGate: A Distributed Intelligent Network Approach to Bridge Switching and Packet Networks, Proc. IEEE ICCCN ’97, Las Vegas, Nevada, September 1997.
- [26] IBM, CallPath Developers Toolkit, 5th ed., June 1996. Available as <http://www.hursley.ibm.com/callpath/khsl4/khsl4mst.html>.
- [27] VocalTec Communications, The VocalTec Telephony Gateway White Paper. Available as http://www.vocaltec.com/products/gtw/gtw_intro.htm.
- [28] C. Low, The Internet telephony red herring, Proc. IEEE GLOBECOM ’96, London, UK, November 1996.
- [29] H. Schulzrinne, Re-engineering the telephone system, Proc. IEEE Singapore Int. Conf. on Networks (SICON) ’97, Singapore, April 1997.
- [30] H. Sinnreich, J. Young, Standards framework for Internet telephony, Draft Proposal for the IETF PINT WG, Munich, Germany, August 12–15, 1997.
- [31] ITU-T, Visual Telephone Systems and Equipment for Local Area Networks which Provide a Non-Guaranteed Quality of Service, Rec. H.323, Geneva, Switzerland, November 1996.
- [32] J. Toga, J. Ott, ITU-T standardization activities for interactive multimedia communications on packet-based networks: H.323 and related recommendations, *Computer Networks* 31 (1999) 205–223, this issue.
- [33] G.A. Thom, H.323: the multimedia communications standard for local area networks, *IEEE Communications Magazine* 34 (12) (1996) 52–56.
- [34] M. Handley, H. Schulzrinne, E. Schooler, SIP: Session Initiation Protocol, Internet Draft, 31 July 1997.
- [35] H. Schulzrinne, J. Rosenberg, Internet telephony: architecture and protocols – an IETF perspective, *Computer Networks* 31 (1999) 237–255, this issue.
- [36] P. Pan, H. Schulzrinne, YESSIR: a simple reservation mechanism for the Internet, IBM Research Report, RC 20967, T.J. Watson Research Center, September 1997.
- [37] W. Almesberger, T. Ferrari, J.Y. Le Boudec, SRP: a scalable resource reservation protocol for the Internet, Technical Report, March 1998. Available from <http://lrcwww.epfl.ch/srp>.
- [38] R. Braden, D. Clark, S. Shenker, Integrated Services in the Internet Architecture: an Overview, IETF Informational RFC 1633, Network Working Group, June 1994.
- [39] L. Zhang, S. Deering, D. Estrin, S. Shenker, D. Zappala, RSVP: a new Resource reSerVation Protocol, *IEEE Network*, 1993.
- [40] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification, IETF RFC 2205, September 1997.
- [41] S. Shenker, C. Partridge, R. Guérin, Specification of Guaranteed Quality of Service, IETF RFC 2212, September 1997.
- [42] J. Wroclawski, Specification of the Controlled-Load Network Element Service, IETF RFC 2211, September 1997.
- [43] H. Schulzrinne, S. Casner, R. Frederik, V. Jacobson, RTP: A Transport Protocol for Real-time Applications, IETF RFC 1889, January 1996.
- [44] K. Nichols, V. Jacobson, L. Zhang, A Two-bit Differentiated Services Architecture for the Internet, Internet Draft <draft-nichols-diff-svc-arch-00.txt>, November 1997.

- [45] R. Guérin, V. Peris, Quality-of-service in packet networks: basic mechanisms and directions, *Computer Networks* 31 (1999) 169–189.
- [46] D. Hutchison, R. El-Marakhy, L. Mathy, A critique of modern Internet protocols: the issue of support for multimedia, *Proc. ECMAST '97*, Lecture Notes in Computer Science, vol. 1242, Springer, Berlin, 1997, pp. 507–522.
- [47] D. Flanagan, *Java in a Nutshell, A Desktop Quick Reference for Java Programmers*, O'Reilly, May 1996.
- [48] H. Kim, F. Steegmans, N. Narayanan, J. Rajahalme, Managing TINA streams that use Internet transport, *Proc. TINA '97*.
- [49] R. Mandeville, D. Shah, Gigabit gets it done, *Data Communications Magazine* (February 1998). Available as http://www.data.com/lab_tests/ethernet.html.
- [50] Avici Systems, *The World of Terabit Switch/Router Technology*, March 1998. Available as http://www.avici.com/html/white_paper.html.
- [51] D. Isenberg, Rise of the Stupid Network, June 1997. Available as <http://www.computertelephony.com/ct/att.html>.
- [52] C. Smith, Applying TINA-C service architecture to the Internet and intranets, *Proc. TINA '97*.
- [53] G. de Zen et al., Value-added Internet: a pragmatic TINA-based path to the Internet and PSTN integration, *Proc. TINA '97*.
- [54] C.A. Licciardi, R. Minerva, C. Moiso, G. Spinelli, M. Spinolo, TINA and the Internet: an evaluation of some scenarios, *Proc. TINA '97*.
- [55] F. Boujema (Ed.), *Introduction of Distributed Computing Middleware in Intelligent Networks*, White Paper P.508. DC-MIN.10, Vers. 10.0, EURESCOM Project P508, 31 July 1997.
- [56] J. Pavón, J. Tomás, Y. Bardout, L. Hauw, CORBA for network and service management in the TINA framework, *IEEE Communications Magazine* 36 (3) (1998) 72–79.
- [57] P. Oechslin, J.-Y. Le Boudec, IMMUNE framework document, Internal Report, Communication Networks Laboratory, Department of Computer Science, Swiss Federal Institute of Technology, Lausanne, Switzerland, 1996. Available from <http://ircwww.epfl.ch>.
- [58] <http://www.tinac.com/>.

Jean-Pierre Hubaux has been an associate professor at EPFL (Swiss Federal Institute of Technology, Lausanne) since 1990, where he is now co-director of the Institute for Computer Communications and Applications. His current area of interest is service engineering, and in particular the convergence of telecom and Internet-based services. At the beginning of his activity at EPFL, he defined the curriculum in Communication Systems; this curriculum integrates EE and CS courses; it is today the division of EPFL attracting the largest number of students. Prior to this position, he spent 10 years in France with Alcatel, where he was involved in R&D activities, mostly in the area of switching systems architecture and software. He was also in charge of the introduction of artificial intelligence techniques, especially for system troubleshooting.

Constant Gbaguidi graduated in electrical engineering at EPFL, Lausanne, Switzerland, in 1995. Since then, he has worked at EPFL as a research and teaching assistant in telecommunication services and management. His diploma thesis was on an implementation of a Virtual Private Network (VPN) over the OSF's Distributed Computing Environment (DCE). In the summer of 1996, he was a Visiting Scholar at Columbia University, New York, where he worked with the COMET group on porting a network prototyping framework to CORBA. Currently, his interests span distributed computing, QoS management, and use of the Internet to facilitate service management.

Shawn Koppenhoefer was born in Canada in December 1968. He successfully completed a double honours program at the Faculty of Mathematics of the University of Waterloo in 1990. From 1990 to 1996 he was a research assistant in the Industrial Computing Laboratory at the Swiss Federal Institute of Technology (EPFL). In 1996, he successfully completed his Ph.D. Dissertation on formal abstractions/models of synchronous time-critical control systems. At the start of 1997 he was invited to join the Telecommunications Laboratory (TCOM/TSG) and he started teaching an introductory course on Networks and Protocols. Currently, he can be found in the Computer Science Department assuming teaching/research/administrative responsibilities in the newly created ICA institute. His professional areas of interest include security in network management, online validation of telecommunication services, cryptographic protocol analysis and the study of formal logics of trust and belief.

Jean-Yves Le Boudec was born in France in 1958. He graduated from Ecole Normale Supérieure de Saint-Cloud, Paris, where he obtained the Aggregation in Mathematics in 1980. He received his doctorate in 1984 from the University of Rennes, France, and became an Assistant Professor at INSA/IRISA, Rennes. In 1987 and 1988, he was with Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department. In 1988, he joined the IBM Zurich Research Laboratory at Rüschlikon, Switzerland, where he was manager of the Customer Premises Network Department. In 1994, he formed the Laboratoire de Réseaux de Communication at EPFL, which in autumn 1997 became part of ICA. His interests are in the architecture and performance of communication systems.