

The Impact of the Internet on Telecommunication Architectures

Jean-Pierre Hubaux¹, Constant Gbaguidi², Shawn Koppenhoefer³ and Jean-Yves Le Boudec⁴

Institute for computer Communications and Applications (ICA), Swiss Federal Institute of Technology, 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 trends, and elaborate on possible ways of salvaging the best parts of the work achieved by the TINA-Consortium in the Internet context.

Keywords: Intelligent Network; Telecommunications Management Network; Telecommunications Information Networking Architecture; PSTN/Internet Interworking; Internet Telephony; Virtual Private Network

1. Introduction

Telecommunications and computer networking are converging. By taking advantages of the potential that each world offers, exciting possibilities are opening new vistas previously unimagined.

In the prehistoric days of telecommunications (late 1970's)⁵, 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. Custom-

-
1. Corresponding author. Email: Jean-Pierre.Hubaux@epfl.ch.
 2. Email: constant.gbaguidi@epfl.ch.
 3. Email: shawn.koppenhoefer@epfl.ch.
 4. Email: jean-yves.leboudec@epfl.ch.

ers 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 heterogeneity, 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, by providing a means to uncouple *service-specific software* from *basic call processing* (this first wave was called the IN - the Intelligent Network [1,2]), and by organizing, across a certain number of layers, the management of telecommunication equipment; this organization facilitates the management of heterogeneous networks (this second wave was called the TMN - the Telecommunications Management Network [3]).

Introduced in the 1980's, the IN was created after the introduction of the 800 number service in the US. Its presence allowed a fast and smooth introduction of new telecommunication services. The publication of its standard followed several years later. The release of the TMN standards occurred roughly at the same time. 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, 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 and addressing the gap between them, a new architecture (TINA - Telecommunications Information Networking Architecture [4]) was created by a worldwide consortium, TINA-C, formed in 1992. The Telecommunications world 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*¹. 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 will generate something new that we are only beginning to discover. The ever-growing popularity of the Internet has dramatically changed the landscape of the communications marketplace. 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 carefully investigate in this paper.

We organize this paper as follows: we briefly recall the basic notions of the various key network architectures introduced to solve the problems perceived by the telecommunication operators. We introduce the IN, TMN and TINA architectures and show their main characteristics, components, interfaces and principles (Section 2). We show how the advantages of the two environments, the PSTN and the Internet, might be used to resolve issues of scalability, service introduction speed, openness, confidentiality, and quality of service. We do this by looking at particular solutions that allow the PSTN to interwork with the Internet, and solutions that replace the PSTN with

5. The term 'prehistoric' is clearly ironic, but our use of this word is intended to reflect the fact that whereas twenty years ago, changes in telecommunication services could be expected to take 5 years at a time, now changes are expected to occur at a rate of 4-5 months. This means that since the start of this decade, we have seen a century-worth of old-time-scale changes. This pace is only accelerating.

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

the Internet (IP Telephony, Section 3). Then we examine a number of assumptions that have changed since the introduction of TINA (Section 4). 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.

2. A Reminder on Network 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. 6.

2.1. The Intelligent Network

The term Intelligent Network (IN) was first introduced by Bellcore in the 1980’s following the deployment of the “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 (ETSI) [5]. Detailed presentations of the IN are provided in [1] and [2].

Before the deployment of the IN, the introduction of new services (such as call forwarding, call screening or alternate billing) in the telephone network required a modification of the software of all the switches in the network. This rendered operators highly dependent on their equipment providers. With the IN, conversely, they could free themselves from this dependency. This is why the network operators have been the main promoters of the IN.

How the IN rendered operators independent is based on some underlying simple principles: 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 themselves are called the *Service Switching Points* (SSP) (Fig. 1). The SSP-to-SCP communication is done using the Common Channel Signalling no. 7 (CCS7) [6].




Fig. 1. Simplified architecture of the Intelligent Network.

Whenever an SSP realizes that a given communication is of a special type (e.g., based on the identification of a special prefix), it will, temporarily, suspend the basic call procedure that it initiated, and will request the collaboration of the SCP.

The development of the IN has been organized in several phases, each of them being characterized by a specific capability set. The oldest one is Capability Set 1 (CS-1), which covers services that are single-ended; a single-ended service is a service that applies to one and only one party in a call and is independent of any other parties that may be participating in a call. Magedanz and Popescu propose to classify the IN CS-1 services into the following categories [2]:

- number translation services (e.g., call forwarding, or universal personal telecommunication)
- alternate billing services (e.g., credit card calling)
- screening services (e.g., originating/terminating call screening, or security screening)
- other services (e.g., malicious call identification, or televoting)

Capability Set 2 (CS-2) [2] contains the following categories of services:

- internetworking services (e.g., internetwork freephone, or internetwork televoting)
- call-party handling services (e.g., completion of call to busy subscriber, or call waiting)
- personal mobility services (e.g., user authentication, or user registration)

Actually, only call-party handling services are fully supported by the CS-2. The next capability set, called CS-3, is claimed to be more a “revolution” than an “evolution”. It includes improvements, such as multiple points of control, IN/ISDN interworking, but also “revolutions” such as broadband and Internet support.

Only the CS-1 is presently wide-deployed; that is why the IN is mostly confined to basic-telephony service control capabilities. This situation may change in the future with the deployment of the CS-3. Much of the discussion in this article considers the CS-1 which was not intended for the Internet.

2.2. The Telecommunications Management Network

In parallel with the IN standardization, the ITU-T has defined the Telecommunications Management Network (TMN) [3]. The TMN helps federate the management of the telecommunication network equipment, generally purchased from different vendors, in order to enable the control of this equipment in a uniform, global and efficient way.

The management of telecommunication networks may be defined as the set of activities of monitoring, analyzing, controlling, and planning the operation of telecommunication network resources in order to provide services with a certain level of quality and cost.

- **Monitoring** is the process of dynamic collection, interpretation and presentation of information concerning the system under scrutiny [7].
- **Analysis** is applied to the information monitored in order to determine average or mean variance values of particular status variables. Analysis is application specific. It can range from very simple computation of statistics to very complex analysis based on sophisticated models [8].
- **Control** is the process of explicitly changing the managed network (topology, parameterization, functionality).
- **Planning** consists in defining the network topology and in dimensioning every network element so that users can get any given service in optimal conditions, in terms of quality and price.

The TMN can be looked at from three points of view: information architecture, functional architecture, and phys-

ical 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 information architecture also defines management layers which correspond to levels where decisions should be made, and where management information resides.

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

The **TMN physical architecture** corresponds to the physical realization of the functional architecture. Each function block becomes a *physical block* or a set of physical blocks, 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). To ensure interoperability, the specification of an interface requires the use of identical communication protocols and identical data representation. The exchange of information between two management systems is performed by means of the Common Management Information Protocol (CMIP).

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 the Private Numbering Plan (PNP), the Closed User Group (CUG), and security [5,9].

Fig. 2 shows the physical architecture of a VPN configuration management system [10,11]. 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 the network provider). The Q3 interface lies between OSs of a given management domain.

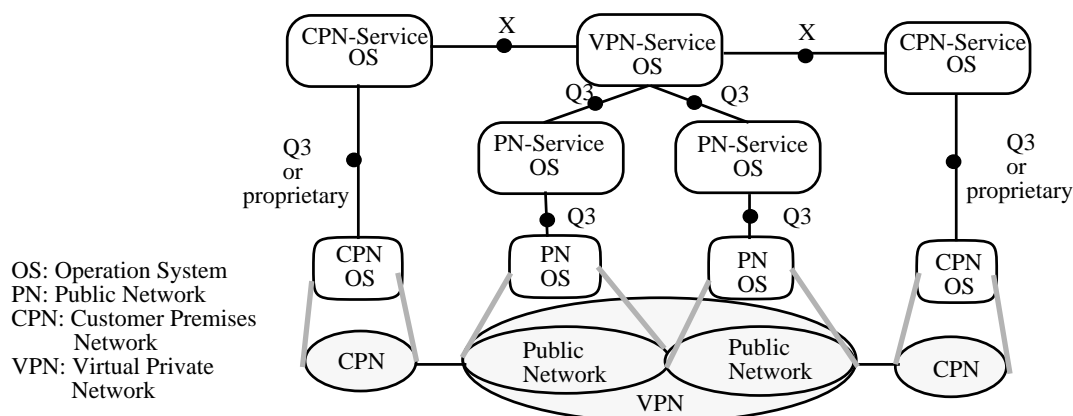


Fig. 2. VPN physical management architecture

The IN and TMN architectures overlap [12]. 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 troublesome. The TINA architecture was expected to solve this complex problem.

2.3. The Telecommunications Information Networking Architecture

In the early 1990's, it appeared that neither the IN, nor the TMN, nor even the coexistence of both, would be enough to respond to the challenge of service provisioning and network management. Few people were convinced that either of those architectures could cope with the complexity of multimedia, multiparty services. Many believed that a new architecture was required, one that could cope with these new issues while still having features belonging to both the IN and TMN domains (e.g., the provision of management-based services). It was felt that the new architecture should also integrate the latest advances in distributed computing (e.g., Open Distributed Processing, ODP [13] and the Common Object Request Broker Architecture, CORBA [14]). The development of distributed software was then perceived to be the most difficult challenge that telecommunication stakeholders had to confront. A worldwide consortium was set up in 1992, named the TINA-Consortium or TINA-C [4].

Fig. 3 sketches out the main principles of TINA. The transmission and switching resources of the network are modeled as *resource components*. These components are used by *service components*, namely *telecommunication*, *information*, and *management services*. Telecommunication services consist in the transport of bits of information between terminals. Information services handle information resources such as movies, sound or documents; they include the storage content, the visualization (the application that is able to interpret the resource), as well as ancillary services such as billing and caching. Management services consist of fault, configuration, accounting, performance and security management, as well as service life-cycle management, service instance management and user life-cycle management. The service components are used to build up TINA services, a TINA service being "a meaningful set of capabilities provided by an existing or intended system to all business roles that utilize it" [15].

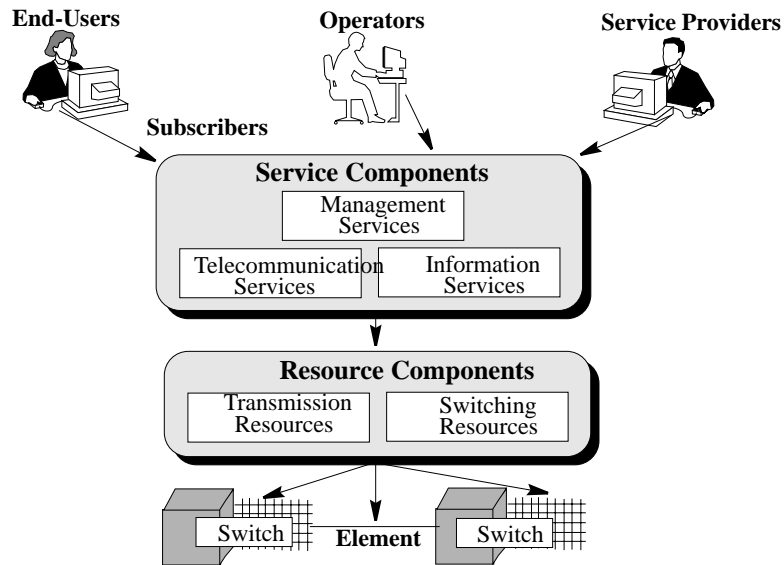



Fig. 3. The TINA architecture

TINA services may be provided according to the business model depicted in Fig. 4. The TINA architecture is not tied to this model, though. The network operator (or more precisely the connectivity provider) sells its resources to the other stakeholders involved in the service architecture. The user, also called the consumer, interacts with third-party service providers in order to obtain the desired services. This interaction is usually achieved through a service retailer. The notion of "service retailer" is related to the expectation that the worldwide service market will be very large; the service retailer helps the user find the service provider he needs. The retailer is also in charge of establishing trust between both parties, notably for security and billing purposes.



↔ Reference points, specified in OMG Interface Definition Language (IDL)

Fig. 4. The TINA business model.

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) [16], which was recently recommended to be the CORBA [17].

In this architecture, a distinction is made between:

- *operational interfaces*, in which interactions are structured in terms of invocations of operations and responses to these invocations, and
- *stream interfaces*, which represent communication endpoints producing or consuming continuous information flows.

- 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 three most important concepts in this architecture are:

- the concept of session, which refers to the service activity;
- the concept of access, which relates to the associations between the user and the service; and
- the concept of management, for service management.

The service architecture (Fig. 5) 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.

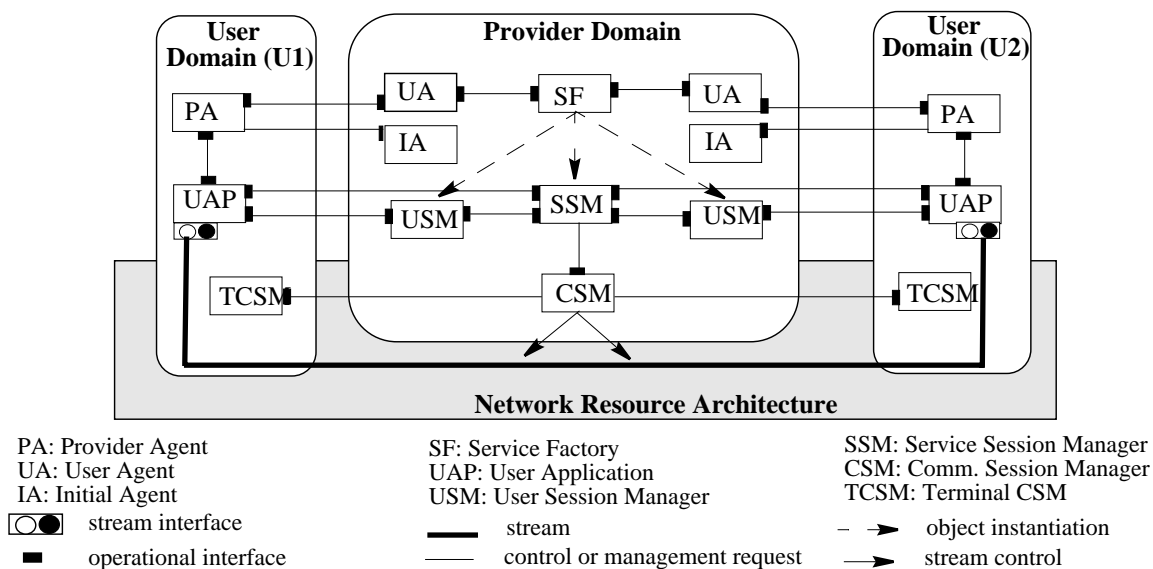


Fig. 5. TINA objects involved in a service scenario

- 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. 6. 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.

The idea of separating, both physically and logically, the supervision of a network from the network itself has been already applied several times in the area of Telecommunications. For example, it can be found in the Common Channel Signalling System Number 7 [6] and in the TMN. In the case of TINA, the reason for this separation is notably that the pace of evolution of the two networks is expected to be very different. Indeed, the evolution of the transport network is tightly constrained by the slow deployment of transmission equipment (including the installation of optical fibers and antennas); so far, this slowness has not been a major competitive weakness, thanks to the monopolistic or oligopolistic situation of network operators. Conversely, the Information Network is directly under the pressure of competition for service provision, and can benefit from the rapid technical progress in the area of distributed computing.

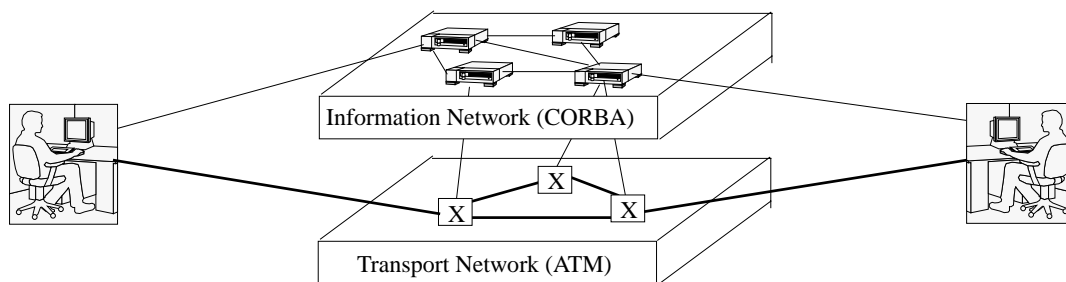


Fig. 6. A simplified view of TINA

Many TINA-compliant connection management implementations have been carried out, especially by the members of the TINA Core Team and by some TINA ancillary projects. A prototype called OAMS (Open management Architecture for Multimedia Services) was implemented by our laboratory in 1994 [18,19]. The main features of our implementation were: connection monitoring, establishment of point-to-point as well as one-to-many connections, and connection release. The network technology considered was ATM. It is worth mentioning that the first implementation of the OAMS architecture was carried out over OSF's DCE (Distributed Computing Environment), which was not object-oriented. We then decided to port the implementation to a more appropriate (i.e., object-oriented) distributed environment, namely CORBA. A future direction for OAMS is the introduction of new management technologies such as intelligent agents, mobile computing, and Web-based management [20].

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 the bridging of 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 [21,22] 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 Network 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 [23]: (1) an open service creation environment for content-based 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) no guarantee of confidentiality, (2) no guarantee of quality of service, (3) the high number and diversity of the Internet Service Providers (ISP), each of whom provides minimal (i.e., best-effort) service and each of whom has very limited motivation to guarantee inter-domain quality of service.

Bringing together the advantages of both worlds would be of mutual benefit by alleviating some of the deficiencies of each of the environments. For instance, the rapidity of service deployment over the Internet could be leveraged when implementing part of the IN architecture, especially the Service Data Point (SDP) which holds information about the user profile and the service logic.

There are *two tendencies* towards the interweaving 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 Internet Engineering Task Force’s (IETF) PSTN/Internet Interworking (PINT) group [24]. Specifically, the PINT group devised a general framework for PSTN/Internet interworking and is currently tackling the specification of the framework components. The PINT contribution is described in section 3.1. Particular realizations that fit into the PINT framework are discussed as well.

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: that is IP telephony (section 3.2). Both tendencies are presented in the following

sections, and some of their strengths and weaknesses will be 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 [25].

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 [23,26] (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 much more *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 came up with 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 [24] (updated in [27]). This ongoing work emphasizes the interfaces between the modules that make up the framework (Fig. 7). The main framework elements identified 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.

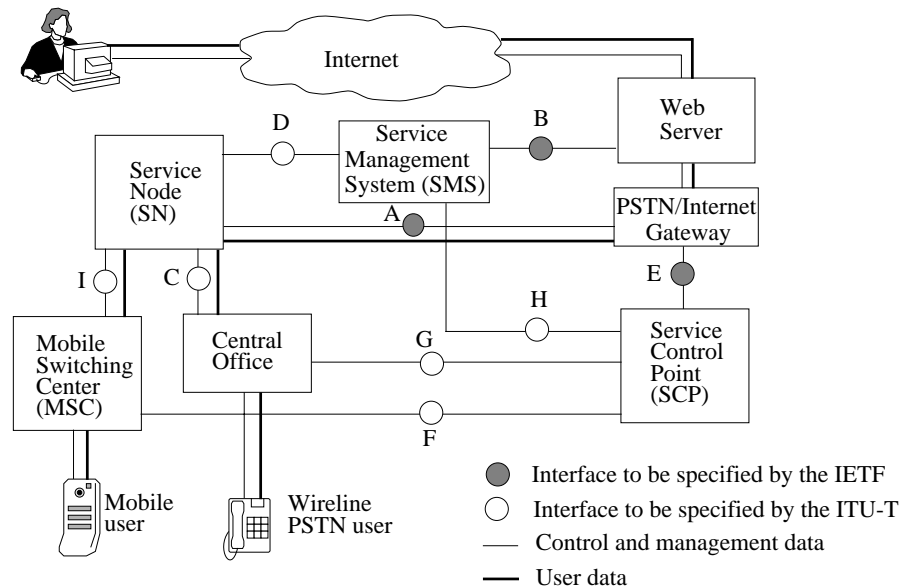


Fig. 7. Reference Model for PSTN/Internet Interworking (PINT) (augmented from [24])

Four services are considered as case studies for the reference model in Fig. 7. These services are: Click-to-Dial, Click-to-Fax, Click-to-Fax-Back, and Voice-Access-to-Content. The Click-to-Dial service enables a user to initiate a PSTN call by hitting a button during a Web session. The Click-to-Fax service is similar to the Click-to-Dial service except that, instead of issuing a request for a voice service, the user clicks on the button to send a fax. The Click-to-Fax-Back service enables the user to request - and eventually receive - a fax by clicking on a button. The Voice-Access-to-Content service enables the user to access content-based services such as Web pages from her telephone set. A Specialized Resource (SR) is used to convert text documents to audio files that are played back to the user.

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 equipment 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. 7). Upon receipt of the request, the SN checks the information supplied against the user profile in order to find out whether the user has subscribed to the Click-to-Dial service. The SN looks up the user profile in order to get the user phone number. The user profile may be updated by the user herself through the Web. In this case, the Web server collects the necessary information from the user and sends it to the SMS through a B interface which is based on the SNMP management protocol [28]. The SN must also determine which sales representative to call; the company may have subscribed to a call distribution (or routing) service. The SN then performs the service logic as indicated by the company's profile. This logic determines the phone number to call and passes it on to the central office - or the mobile switching center - to initiate the call. This is done through either the C or I interface.

In the Click-to-Dial example, the SCP may be used as well, in lieu of the SN. However, this may not be the case for the Click-to-Fax service which needs the conversion of the fax document, sent by the user, into a fax format acceptable by the PSTN. Indeed, SCPs do not hold the specialized resources needed for conversion operations.

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 [29].

Two protocols are competing for the interface A between the PSTN/Internet gateway and the SN: the Service Support Transfer Protocol (SSTP) [30] and the Session Initiation Protocol (SIP) [31]. SSTP is a protocol designed for the specific needs of the PINT reference model. SIP, on the other hand, is a generic protocol intended to enable invitation to users to participate in multimedia sessions. Additionally, SIP enables user mobility by forwarding invitations to the user's current location. Schulzrinne made use of the SIP in his proposition of a control system for the Click-to-Dial-Back service [32]. Being devised to suit the specific needs of PSTN/Internet interworking, SSTP is obviously much simpler and likely to be more efficient than SIP.

Many implementation efforts were carried out in the area of Internet/PSTN interworking. In the following section, we report on some of the projects that we could find in the literature.

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 [33], Hewlett-Packard [26], IBM [34], and Vocaltec Commu-

nications [35]. None of their solutions, except that from Hewlett-Packard, addressed the particular issue of porting the IN architecture to the Internet. Most projects tackled only part of the PINT reference model. Consequently, we concentrate on the Nexus project from the HP laboratories in Bristol, UK. Some of the other aforementioned efforts fall better into the IP telephony category, since they deal with continuous data flow through Local Area Networks (LAN). Therefore, they are discussed in section 3.2.

Nexus was initiated with the objective of integrating the Internet and Telephony in order to create open, scalable, and integrated services [23,26,36]. The Web already provides for such an easy service creation environment. The challenge is to bring this ease of development to the IN.

The main outcome of the Nexus project is a prototype architecture called WebIN (Fig. 8). 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 might be 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. 8 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.

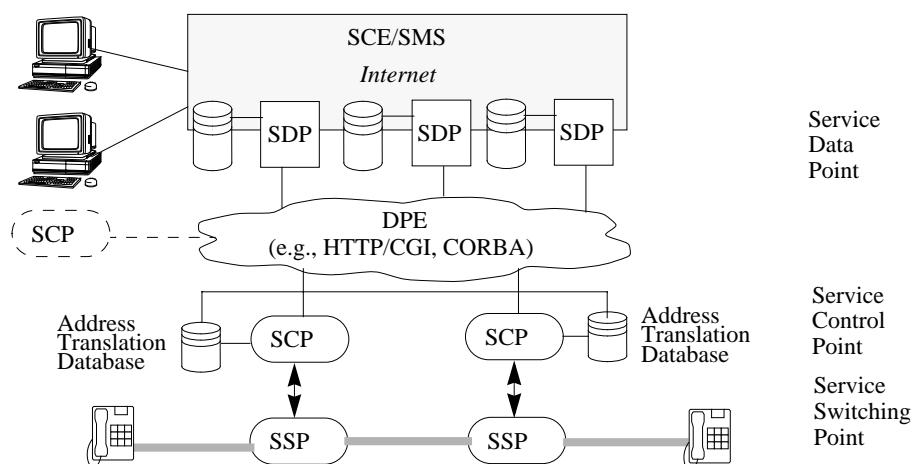


Fig. 8. The WebIN physical architecture (adapted from [36])

3.1.3 Provision of the VPN Service

In this section, we show how well a fairly complex service such as the VPN is implemented 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 as many 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. 9). 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. Therefore, the VPN data has to be duplicated as many times as the number of IN operators involved.

An important shortcoming of this solution is that changes to the VPN must be notified to each and every IN operator implicated in the provision of the overall service. In the very likely situation where all the databases are not updated at the same time, a number of miscomputations may appear. 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 from the two CUGs for some time.

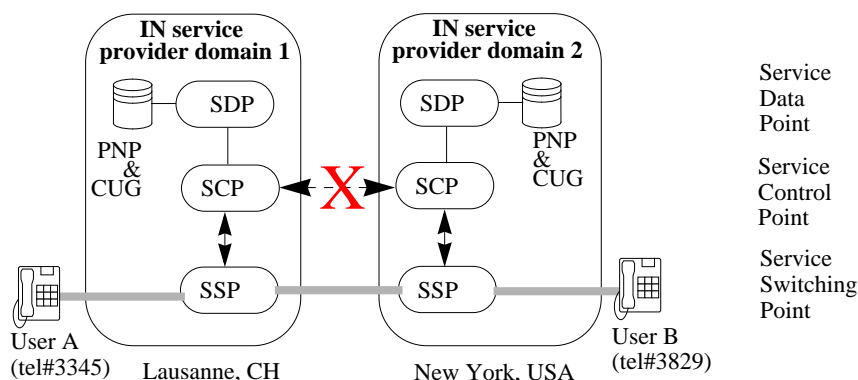


Fig. 9. Provision of a VPN service across two IN provider domains

3.1.3.2 Running the VPN Control System on the Internet

Fig. 10 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 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. Changes in the VPN profile are immediately reflected. Subsequently, the lack of *interoperability* between legacy 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 any formal interaction with the VPN service providers.

To facilitate information search across a network of servers, a spanning tree covering these servers may be used. This spanning tree may be employed to broadcast search information. Likewise, a DNS-like server may be dedicated to handle the location of service data across the server network. For instance, if the data related to a given CUG is needed, then a request for location is sent to the DNS-like entity, which looks up in its database in order to find out the correct location of the data requested.

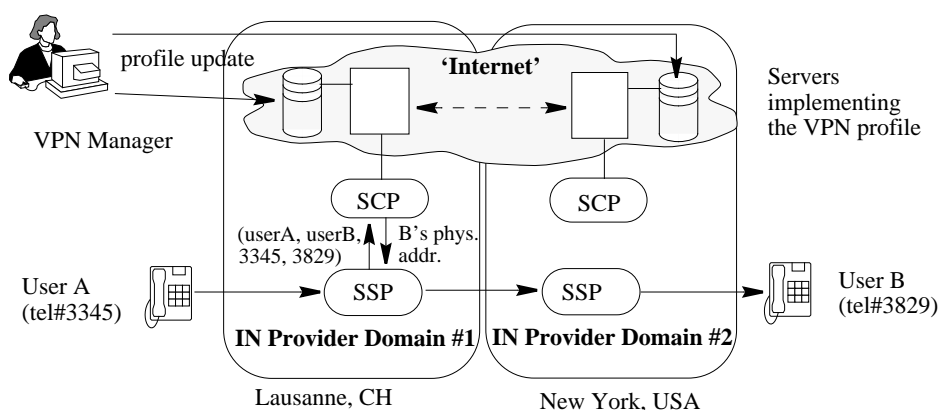


Fig. 10. A realization of the VPN service with the control system running on the Internet

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

Schulzrinne [38] gave an outstanding assessment of the situation of both the Internet and the legacy telephone system. Among the motivations for re-engineering the latter, in the light of the growing dominance of the Internet, he stressed:

- the choice of the compression scheme by the user: the data compression ratio is understood to determine the cost of the communication; by being able to select the compression scheme, the user can control the cost
- silence suppression: in packet-switched networks, silence is naturally suppressed, since no packet has to be processed during this period of time
- traffic separation: fax traffic, for instance, cannot be discriminated from conversation traffic in existing PSTNs. This can be achieved in the Internet by differentiating packets
- ease of service creation on the Internet
- the lower cost of non-technical factors (e.g., the huge amount of administration personnel involved in operating telecommunications largely accounts for the high price of telephony services.)

There are a number of 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. All these issues, especially the last one, are being worked on intensely, as we report below.

In section 3.2.1, the overall framework for IP telephony, proposed by Sinnreich and Young, is described. This framework is, so far, one of the most elaborate found in the literature. In section 3.2.2, we propose a technical implementation framework that identifies the main protocols to be used with providing IP telephony. In section 3.2.3, the provision of the VPN service within the IP telephony framework is briefly discussed.

3.2.1 Architecture

A standard framework for IP telephony (Fig. 11) was proposed by Sinnreich and Young in [39]. 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.

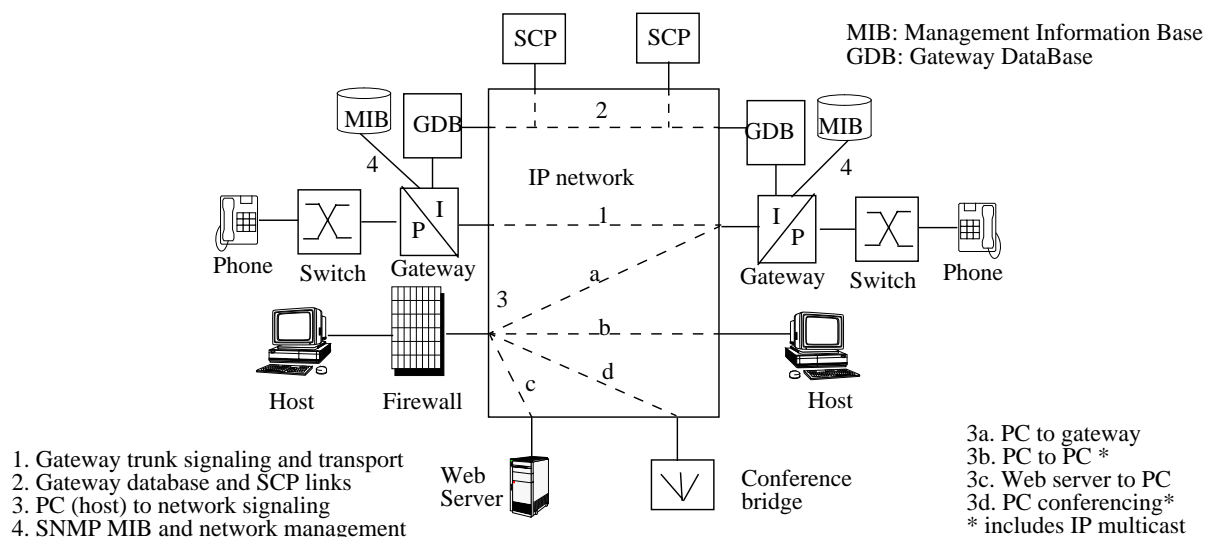


Fig. 11. A standards framework for IP telephony [39]

3.2.2 A Technical Implementation Framework

To find a technical implementation framework, let us examine each relevant element in Fig. 11, and devise its implementation in the context of IP telephony.

The elements in Fig. 11 that need attention are: the host, the Internet/PSTN gateway, and the IP network. Each element is considered in turn in the next sections. The technical implementation framework proposed below is more complete and up-to-date than the one in [39]. Specifically, we cover interesting contributions such as the SIP and

the Windows Telephony Application Programming Interface (TAPI) [40].

3.2.2.1 The Host

We distinguish and discuss two issues in this section: telephony session control, and connection control. We use “session” here in the sense of a conference, i.e., an aggregation of interacting users; session control is concerned with inviting a new user to join a session, resolving user location, adding a new medium, and so forth. Connection control refers to the control of the connectivity service supplied by the network.

Session control - The invitation to a user to join a telephony session can be performed using either the Session Initiation Protocol (SIP) or H.323 [41,42]. SIP is a protocol that enables the invitation of users to participate in multimedia sessions. H.323 addresses the provision and control of audio/video interactive services over LANs with/without Quality of Service (QoS) guarantees.

SIP can be used in conjunction to H.323. In this mode, it essentially serves to locate the user; the actual connection is set up and controlled according to H.323. Handley *et al.* suggested the use of SIP with Internet multimedia conferencing [43]. Although, H.323 and SIP share some common functionality, the former is not specifically intended for the Internet, even though it might be easily adapted to the Internet [44]. Moreover, SIP is an ongoing work which is being enhanced with new call control services [45]. This explains the growing attraction for SIP as the protocol to use before the connection is setup.

Connection control - SIP does not address connection control. We can therefore look to H.323. Note that the trend toward using H.323 for connection control can be confirmed by the newer version (3.0) of the TAPI library from Microsoft which already takes into account H.323, as well as QoS and multicasting concerns [40].

The Real-time Transport Protocol (RTP) [46] will be used - as it is included in H.323 - by applications to transfer real-time information over the Internet. Its accompanying control protocol, the Real Time Control Protocol (RTCP), will be used to control real-time communication sessions on an end-to-end basis. The Resource ReReservation Protocol (RSVP), implemented at the host, would be used for reserving and managing the communication resources [47,48]. Note, however, that its complexity has steered new alternatives, such as:

- the YEt another Sender Session Internet Reservation (YESSIR) protocol [49], which is an extension of RTP to support resource reservation along the path between the sender and the receiver
- differentiated services, which implement resource reservation essentially by tuning scheduling policies [50]. Packets traveling on the Internet are marked according to the QoS expected by the user. When congestion is foreseen by the router, packets receive different treatments depending on the type of service that they represent. Hence, packets relating to the traditional best-effort service would be dropped first when congestion is imminent.

The main strength of these alternatives over RSVP is that they demand little modification in the routers as currently implemented.

As we can notice, separating session control from connection control helps discuss each concern individually, instead of just comparing SIP to H.323. We believe that these two approaches will co-exist in the future.

3.2.2.2 *The Internet/PSTN gateway*

In the previous sections, we used the term “gateway” with different meanings. The Internet/PSTN gateway of the PINT reference model (Fig. 7) does not carry *continuous* data flow. It is mainly designed to forward the user requests from the Internet domain to the PSTN. In some special cases, such as Click-to-Fax, the PINT gateway can carry the user data, but never on a continuous basis.

In this section, we use the term “gateway” to refer to an application level gateway (since the PSTN and Internet domains use completely different protocol stacks) that carries continuous data flow. A generic description of such a gateway is given by Simeonov *et al.* [33]. The main elements of this gateway are: (1) the service node controller, which allocates and releases the resources needed by each service; (2) the resource platforms, which implement functions such as dedicated media processing (e.g., voice/fax reception), and special resource functions (e.g., collecting user input, and announcement playout); (3) the channel matrix switch, which routes calls from the PSTN to the appropriate resource platforms; (4) the data packet switch, which routes calls from the Internet to the appropriate gateway area; (5) media conversion processors, which convert formats across the PSTN and Internet domains; and (6) the Internet gateway, which essentially implements the bus on which data is conveyed from the Internet to the PSTN, and vice versa.

A description of gateways from the protocol standpoint can be found in [41] as part of the H.323 framework.

3.2.2.3 *The IP Network*

Guérin discusses the provision of QoS in the Internet in this special issue [51]. We simply intend to recall the main trends towards satisfactory provision of telephony services on the Internet. We consider two standpoints for the following discussion: session control and connection control.

Session control - As mentioned in section 3.2.2.1 for the host, either SIP or H.323 can be used to control telephony sessions within the IP network. Schulzrinne discusses the use of SIP for telephony services in this special issue [52].

Connection control - To be consistent with the design of the host, we suggest the use of H.323 to control the connectivity service offered by the Internet. The reservation style may be any of that by the IETF Integrated Services (int-serv) group, the one proposed by the Differentiated Services (diff-serv) group, or YESSIR. The int-serv group proposes an integrated services architecture [53] that recommends RSVP as the reservation protocol. Two network services can currently be requested with RSVP: guaranteed service [54] and controlled-load service [55]. 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.

In spite of these advances, the IP paradigm for real-time services is still inconsistent in many respects. A relevant critique of this paradigm is given in [56].

3.2.3 *Provision of the VPN service*

The provision of the VPN service within the IP telephony realm is straightforward from Fig. 10: 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.

It is worth mentioning that there is another so-called VPN service already implemented that might cause confusion with the one we have been discussing here. The VPN, as we view it, is based on the PSTN. On the other hand, many commercial products claim to provide VPN services over the Internet. Examples are Microsoft's VPN service based on Windows NT [57], and Data Fellows' F-Secure VPN [58]. These products provide the users with access to their home Intranet from anywhere at any time in a secure way. The main technology underlying Microsoft's solution is a tunneling protocol called PPTP (Point-to-Point Tunneling Protocol). PPTP can encapsulate many network layer protocols (IPX, IP, etc.) and carry them over IP networks with a reported high level of security. PPTP can interoperate with PPP (Point-to-Point Protocol) [59], another tunneling protocol basically intended to cope with the diversity of network layer protocols. The VPN service furnished by Microsoft looks much like a secure and mobility-enhanced dial-up service, since users can access the environment of their home (or company) site from anywhere. The F-Secure VPN product from Data Fellows is based on TCP/IP and does not make use of any specific tunneling protocol. The main feature of F-Secure VPN is data encryption.

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 reported on the state-of-the-art work, and assessed how well the works found in the literature support a fairly complex service, i.e., the VPN. We could show that interoperation among IN infrastructures can be easily achieved when the Internet is used to run the service control system (Fig. 10). Moreover, the customer can update the profile of his VPN (by changing the service logic on a Web server) 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 [56]. 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

As we have seen in Section 2, the main motivation behind TINA was to define an architecture that would be able to cope with new services such as multimedia, multipoint communications and their management. 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)
- well-established pricing and billing mechanisms
- 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. The GSM required the development of specific and highly sophisticated hardware, unlike most of the services targeted by TINA.

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 defined and partially deployed at the time TINA was launched; the IN had proved its ability to encompass 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 the latent power of the Internet to support the multitude of capabilities we have witnessed since then.
- The network operators had an understandable reluctance to see their role reduced to that of humble connectivity providers.

In the framework of TINA, this assumption naturally led to a high level of complexity in the Service Architecture.

In the meantime, the Internet services became extremely popular; the pervasiveness of browsers, the fast deployment of Internet telephony and the multiparty sessions over the Mbone showed that it was possible to implement a high number of services over a network of limited capabilities. Indeed, the role of the terminal became completely different; the main reasons for this are:

- the processing power made available with the current technology makes it possible to implement sophisticated functions in software; this includes not only audio and video decoders, but the more complex encoders as well;
- the network itself can be used today by millions of users to download the code of new applications or new versions of applications in virtually no time. The downloading of software can even be set up to occur without the end-user being aware of it; the main motivation behind the Intelligent Network (deployment of new services in months rather than in years) sounds rather weak compared with this impressive possibility;
- software paradigms which recently received wide acceptance thanks to the Java language [60] (applets, Java Virtual Machine) even make it possible to download executable code *during* a given service session.

The TINA Consortium devoted a lot of attention to the definition of complex call models of the telecommunication services (see section 2.3). This is certainly a stimulating challenge, but we think that the most important business opportunities are, and will be, in the area of the services involving one or more computers (belonging to the service provider or to the network operator). As we have seen, these services are called *information services* in TINA parlance; they are often called *Value Added Services* in other frameworks. They clearly encompass content services (e.g., video on demand, and travel reservations) as well as a multitude of computing/networking facilities that businesses (especially small ones) will be willing to outsource. These facilities are for example Web caching, electronic mail, newsgroups, and electronic commerce. The deployment of these services on a large scale and in a coherent way is a major technical challenge, which TINA addressed without considering that all these services would be Internet-based; for this reason, the Service Architecture [15] remains at an inappropriately high level of abstraction.

The bottom line is that network operators and service providers should concentrate on the provision of services (or service mechanisms) which *cannot* reasonably be supported only by the terminals. Many value-added services present this characteristic; additional 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, take advantage of smart terminals (e.g., Network Computers, or television browsers), 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.

This assumption explains why so much emphasis has been laid on connection management in the framework of TINA [4]. The possibility of coping with a connectionless transport network has only been taken into account recently [61].

It is now clear that the ITU-T's old vision of a homogeneous, ubiquitous, general-purpose ATM network supporting the B-ISDN has vanished. New technologies, such as the Gigabit Ethernet [62] and Terabit Switches/Routers [63], can deliver satisfying QoS at a lower cost than ATM.

Certainly, ATM still has a future as a backbone technology, or in combined scenarios with other protocols, such as IP over ATM [64], multiprotocol encapsulation over ATM [65], and AREQUIPA [66]. However, local area networks will continue to be primarily frame- rather than cell-based; hence, the end-systems will continue to have to cope with this technology. It becomes more and more likely that the B-ISDN will in fact be based on the Internet/Intranet technologies, provided that the appropriate resource reservation, billing and security mechanisms are implemented. Therefore, any general-purpose service architecture must be able to cope with both connection-oriented and connectionless networks.

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*.

Significantly, the TINA Core Team was set up for a duration (1992 - 1997) roughly equivalent to a study period. It was still believed that the evolution was under control; devoting *five years* to the design of an architecture, to the definition of the key interfaces and to the implementation of some large-scale prototypes was thought to be a reasonable goal.

For the same reason, advertising the results of TINA was not considered to be a top priority for the consortium. On the contrary, several of the documents were kept *confidential* until 1997, and the whole contribution has remained somewhat esoteric because of its abstractness and terminology.

Meanwhile, the development of the Internet protocols has proved that it is possible to "do first and standardize later". *Incremental development* is now the usual practice. This means that some traditional ideas of software engineering such as the waterfall model are questionable when applied to protocol development. With this in mind, it is easy to see that TINA is suffering from the same disease as the one that nearly wiped out the OSI protocols.

TINA is not the only victim of the change introduced by the Internet; the legacy telephone network was also shaken out. Isenberg [67] assessed this impact by analyzing, as we did for TINA, the assumptions made by operators of the telephone network. Mainly, these operators believed that voice traffic would be preponderant over other types of traffic in the telephone network, and that the user telephone set would remain dumb. We all know today that modems are widely used to exchange data over the PSTN, and that the users can have more control on their communications (e.g., choice of the encoding technique) when they hold a computer.

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 [56]. 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.

Smith [68] identifies a number of problems related to the Internet services, both for the developers and the users. He argues that the *development* of interactive services is a painful process. This is explained by the fact that support-services have to be hand-built for each service, and that there is a lack of an appropriate paradigm for the communication between the browser and the server. Likewise, the *use* of Internet services is claimed to be uncomfortable, because these services are obtained through many points of contact, each with their own authentication system and payment mechanism. In both cases, Smith argues that these problems can be solved by an appropriate use of specific TINA concepts: the separation of service provision and service retailing, the generic session model and the CORBA technology.

De Zen *et al.* [69] present the SISTINA architecture. This very pragmatic approach provides an integrated supervision of connectionless and connection-oriented communication sessions. It is claimed to be easily deployable, because it relies on the existing network infrastructure.

Licciardi *et al.* [70] study several scenarios in which TINA could be applied usefully to the Internet. These scenarios cover the management of routers, the connectivity with QoS in the Internet, the TINA service access by means

of Web browsers, and the provision of services over mixed packet- and circuit-switched networks. The areas in which TINA is deemed to have the best chances for successful application is the management of IP over ATM networks.

In [71], representatives of several European operators explain how useful it would be to introduce object-oriented middleware in order to enhance the existing Intelligent Network; CORBA is their favored solution. They claim that this would pave the way to the subsequent introduction of TINA. They identify management as the area in which CORBA is likely to find its earliest acceptance. However, they also stress that the introduction of CORBA in the Intelligent Network is a challenging task. Two possible migration scenarios are proposed: (1) interworking between Signalling System 7 and CORBA, by means of an appropriate gateway; and (2) using Signalling System 7 for the communication between CORBA nodes.

What can be salvaged from TINA? Certainly, the idea of deploying an Object Request Broker (ORB), such as CORBA, in order to support sophisticated supervision functions, has a strong potential. However, it is clear that the Service Architecture presented in section 2 does not make the cut. It is, in many respects, too heavy (both in terms of complexity and performance) to compete with solutions such as the combined use of RSVP and RTP (it is true that these IETF proposals do have some weaknesses, but ongoing experiments hope to resolve these issues). It is not realistic 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 a 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, through the TINA Management Architecture.

A possible integration of the Internet with the TINA architecture, based on the previous analysis, is shown in Fig. 12 (which is derived from Fig. 6 by replacing ATM with IP). Unlike with TINA, the establishment and release of short-lived flows, or connections, in the architecture we propose here would be fully achieved within the Transport Network itself. This would resolve the problem of the lack of performance observed when the connections are entrusted to the Information Network (due to the high latency of CORBA). In this way, the protocols deployed in the Transport Network would be upgraded every time some progress is made in the Internet protocols (e.g., new versions of routing protocols). 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 might be based on a distributed object-oriented framework such as CORBA. This evolution is compliant with the trends in network management research [72]. Communications among the nodes of the Information Network and the routers of the Transport Network would be supported by SNMP.

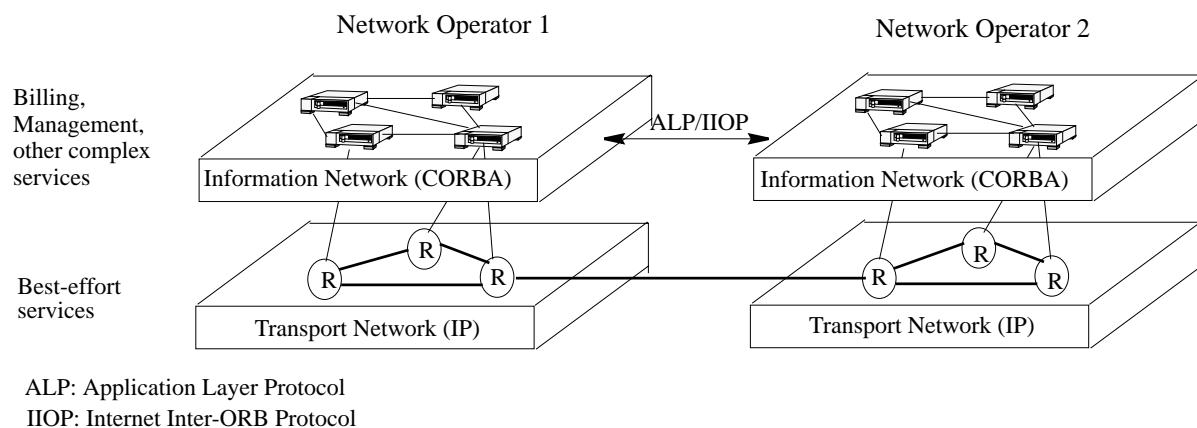


Fig. 12. A possible integration of the Internet with the TINA architecture

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 CMIP in the case of the OSI management model. The interface between two cooperating 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 (IIOP) supports the communication between kTNs of different network operators, and may help implement the functionality of a TMN X interface (Fig. 12).

This approach, however, does not fully meet the needs of the customers. Indeed, there is a strong demand for services providing an assured QoS. Hereafter, *we define 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 [54]) or looser (e.g., the controlled-load service [55]). In particular, many business companies would like to interconnect their distributed Customer Premises Networks (CPNs). Assured services are key to providing seamless Intranet connectivity.

In order to satisfy the need for assured services, 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 provision of this connectivity could be achieved 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 upcoming concepts such as bandwidth brokers and differentiated services [50]; in an Intranet, over-provisioning transmission and routing resources can be a primitive - but effective - way to guarantee QoS. Finally, these assured-QoS connections could also be provided in some cases by *circuit switching*, for example by means of the SDH, for long-lived connections.

The services belonging to the assured service category compete for the same network resources as the ones used by the best-effort services. The Information Network will split the resources between these two categories of services according to their needs; a way to optimize this resource allocation is described in [73]. The Information Network is also in charge of establishing and releasing long-lived connections with assured QoS, while the Transport Network is concerned with establishing and releasing short-lived connections using RSVP; the longer time setup is the price that the end-user must accept to pay in order to get assured QoS. Obviously, the users will be charged

for the resources reserved for them. Fig. 13 illustrates the way the integration of assured and best-effort services can be achieved. This is a simple extension of Fig. 12, where we separate the network resources to support the two service categories concerned with the integration.

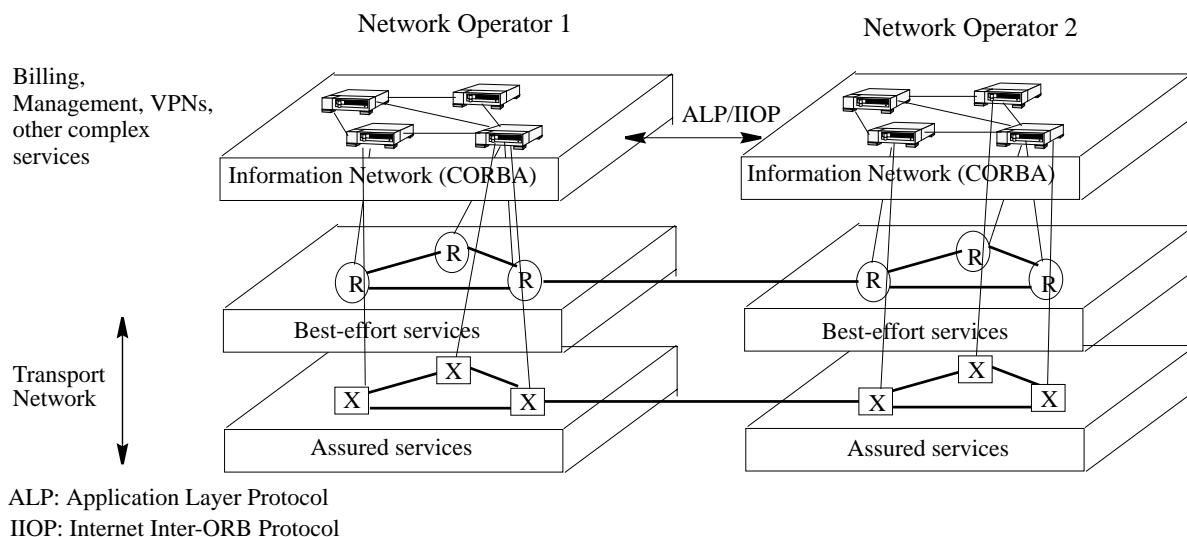


Fig. 13. Integration of best-effort services and services with reservations

If CORBA proves to be successful, one could expect the general architecture for the provision and management of services to look like the one proposed in Fig. 14. The interfaces between application pieces, or service components, are specified using the Interface Definition Language (IDL). CORBA's deployment is already foreseen in the near future as a support for distributed network management platforms [72]. It will however be subject to the fierce competition of other solutions such as the Web, Java Remote Method Invocation (RMI) and the Distributed Component Object Model (DCOM) [74].

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.

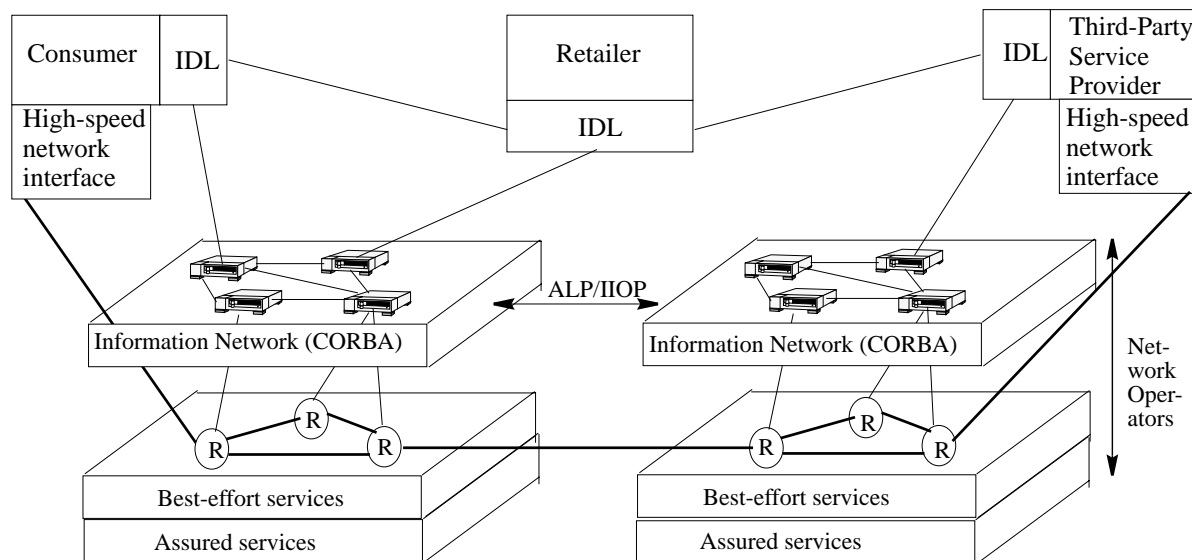


Fig. 14. Service provision and management based on CORBA

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.

To summarize, there are ways to salvage important TINA results and to combine the best of the Internet and TINA worlds. By this means, the user will get both the flexibility (and evolution) features of the Internet, and the dependability of the telecommunication networks.

This requires, however, a radical change of philosophy of the TINA initiative:

- the 3 assumptions discussed in Section 4.1 must be reconsidered
- the importance of the Internet must be recognized
- the focus must be on the most important technical challenge: the provision of information services based on the Internet (telecommunication services can be provided by an architecture like the one shown in Fig. 13; little differentiation, and limited profit, is to be expected from this area)
- the attempt to define an abstract framework from scratch must be avoided, even for management services.

The end of the first phase (1992-1997) and the disbanding of the TINA Core Team are an opportunity for this substantial revision of the objectives. The second phase of the TINA project (1998-2000) [75] now integrates part of the shift in philosophy that we have mentioned above. The importance of the Internet is being partially taken into account, and more emphasis is being laid on the implementation of the solutions proposed by the working groups; these groups will be technically leading the TINA project, like similar groups do at the IETF.

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 trend 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 trend 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 trend towards using the Internet to deploy telephony services, promotes the Internet as an integral transport infrastructure: this is IP telephony. We described the way IP telephony may be provided today, by integrating the main recent advances at the most important levels involved, namely the host, IP network, and PSTN/Internet gateway.

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 the telecommunications community. 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.

It would be an interesting endeavor to deepen the study of 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 will investigate more closely.

Acknowledgments

The authors are indebted to the Guest Editors and the anonymous reviewers for their relevant feedback which help 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, 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 and N. J. Shah, *The Intelligent Network Standards, their Application to Services* (McGraw-Hill, 1996).
- [2] T. Magedanz and 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 and Y. Inoue, The TINA Initiative, *IEEE Communications Magazine*, (March 1993) pp. 70-76.
- [5] ITU-T, Introduction to Intelligent Network Capability Set 1, Rec. Q.1211, March 1993.
- [6] A. R. Modaressi and R. A. Skoog, Signaling System No. 7: A Tutorial, *IEEE Communications Magazine*, (July 1990), pp. 19-35.
- [7] J. Joyce *et al.*, Monitoring Distributed Systems. *ACM Trans. on Comput. Syst.*, Vol. 5, No. 2, May 1987.
- [8] M. Mansouri-Samani and M. Sloman, Monitoring Distributed Systems (A Survey), Technical Report, Imperial College, Dept of Computing, April 1993.
- [9] 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.
- [10] J.-P. Gaspoz, J.-P. Hubaux and S. Znaty, A Generic Architecture for VPN Configuration Management, *Journal of Network and Systems Management*, Vol. 4, No. 4, Dec. 1996, Plenum Press, pp. 375-395.
- [11] J.-P. Gaspoz, Methodology for the Development of Distributed Telecommunication Services. *Journal of Systems and Software*, 1996.
- [12] G. Pavlou and D. Griffin, An Evolutionary Approach towards the Future of IN and TMN, *ICN Special Issue on Telecommunication Services*, Baltzer Science Publisher, Jan. 1998.
- [13] ITU-T, Basic Reference Model of Open Distributed Processing - Part1: Overview and Guide to Use, Rec. X901, 1994.
- [14] OMG, The Common Object Request Broker: Architecture and Specification, Rev. 2.0, July 1995.
- [15] TINA-C, Service Architecture, Vers. 5.0, 16 June 1997.
- [16] TINA-C, TINA Distributed Processing Environment (TINA-DPE), TINA-C Deliverable No. TB_PL001_1.0_95, July 1995.
- [17] S. Vinoski, CORBA: Integrating Diverse Applications within Distributed Heterogeneous Environments, *IEEE Communications Magazine*, Feb. 1997.
- [18] S. Znaty, Service and Network Management in the OAMS Open Service Architecture, *ACM Computer Communication Review*, Vol. 26, No. 3, pp. 59-74, July 1996.
- [19] S. Znaty, F. Torasso and J.-P. Hubaux, Network Management in the OAMS Open Service Architecture, in: F. A. Aagesen, H. Botnevik, and D. Khakhar, eds., *Information Network and Data Communication*, (Chapman and Hall, 1996), pp. 246-59.
- [20] J.-P. Martin-Flatin and S. Znaty, Annotated Typology of Distributed Network Management Paradigms, SSC Technical Report SSC/97/008, Communication Systems Division, Swiss Federal Institute of Technology, Lausanne, Switzerland, 14 March 1997. Available from <http://sscwww.epfl.ch/>.
- [21] P.-A. Etique, J.-P. Hubaux and X. Logean, Service Specification and Validation for the Intelligent Network, *Baltzer ICN, Special Issue on Telecommunication Services*, Jan. 1998.
- [22] 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.
- [23] J. O'Connell and P. Hoft, Web-IN: Integrating Intelligent Networks and the Internet, HP white paper, 30 Nov. 1996.

- [24] I. Faynberg, M. Krishnaswamy and H. Lu, A proposal for Internet and Public Switched Telephone Networks (PSTN) Internetworking, Internet Draft, March 1997.
- [25] RACE Project PREPARE, IBC VPN Services, Specification CFS D721, Dec. 1993.
- [26] C. Low, Integrating Communication Services, *IEEE Communications Magazine*, 35(6), June 1997.
- [27] M. Krishnaswamy, PSTN-Internet Interworking- An architecture overview, Internet Draft, Nov. 1997.
- [28] J. Case *et al.*, Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMPv2), IETF RFC 1905, Draft Standard, Jan. 1996.
- [29] L. Conroy, M. Lautenbacher and 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.
- [30] M. Krishnaswamy and H. Lu, Information exchange to be supported by the service support transfer protocol (SSTP), Internet Draft, July 1997, work in progress.
- [31] M. Handley, H. Schulzrinne and E. Schooler, SIP: Session Initiation Protocol, Internet Draft, 31 July 1997.
- [32] H. Schulzrinne, SIP for Click-to-Dial-Back and Third-Party Control, Internet Draft, 21 Nov. 1997.
- [33] P. Simeonov *et al.*, @INGate: A Distributed Intelligent Network Approach to Bridge Switching and Packet Networks, in: *Proc. IEEE ICCCN '97*, Las Vegas, Nevada, Sept. 1997.
- [34] IBM, CallPath Developers Toolkit, 5th. edition, June 1996. Available as <http://www.hursley.ibm.com/callpath/khsl4/khsl4mst.html>.
- [35] VocalTec Communications, Inc., The VocalTec Telephony Gateway White Paper, available as http://www.vocaltec.com/products/gtw/gtw_intro.htm.
- [36] C. Low, The Internet Telephony Red Herring, in: *Proc. IEEE GLOBECOM '96*, Nov. 1996, London, UK.
- [37] C. Low, D. Skov and N. Raguideau, WebIN- An Architecture for Fast Deployment of IN-based Personal Services, in: *Proc. IEEE Intelligent Network Workshop, IN'96*, April 1996, Melbourne, Australia. Available from <http://callculus.hpl.hp.com:8000/webtel/papers.html>.
- [38] H. Schulzrinne, Re-engineering the Telephone System, in: *Proc. IEEE Singapore International Conference on Networks, (SICON) '97*, Singapore, April 1997.
- [39] H. Sinnreich and J. Young, Standards Framework for Internet Telephony, Draft Proposal for the IETF PINT WG, Munich, Germany, Aug. 12-15, 1997.
- [40] Microsoft Corporation, IP Telephony with TAPI 3.0 White Paper, 1997.
- [41] ITU-T, Visual Telephone Systems and Equipment for Local Area Networks which Provide a Non-Guaranteed Quality of Service, Rec. H.323, Nov. 1996, Geneva, Switzerland.
- [42] G. A. Thom, H.323: The Multimedia Communications Standard for Local Area Networks, *IEEE Communications Magazine*, 34(12), Dec. 1996, pp. 52-6.
- [43] M. Handley, J. Crowcroft, C. Bormann and J. Ott, The Internet Multimedia Conferencing Architecture, Internet Draft draft-ietf-mmusic-confarch-01.txt, July 1997.
- [44] P. Lantz, Usage of H.323 on the Internet, Internet Draft, draft-rfced-info-lantz-00.txt, Sept. 1997.
- [45] H. Schulzrinne and J. Rosenberg, SIP Call Control Services, Internet Draft draft-draft-ietf-mmusic-sip-cc-00.ps, March 1998. See also <http://www.cs.columbia.edu/~hgs/sip/>.
- [46] H. Schulzrinne, S. Casner, R. Frederik and Van Jacobson, RTP: a transport protocol for real-time applications, IETF RFC 1889, Jan. 1996.
- [47] L. Zhang, S. Deering, D. Estrin, S. Shenker and D. Zappala, RSVP: A New Resource ReSerVation Protocol, *IEEE Network*, 1993.
- [48] R. Braden, L. Zhang, S. Berson, S. Herzog and S. Jamin, Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification, IETF RFC 2205, Sept. 1997.
- [49] P. Pan and H. Schulzrinne, YESSIR: A Simple Reservation Mechanism for the Internet, IBM Research Report, RC 20967, T. J. Watson Research Center, Sept. 1997.
- [50] K. Nichols, V. Jacobson and L. Zhang, A Two-bit Differentiated Services Architecture for the Internet, Internet Draft <draft-nichols-diff-svc-arch-00.txt>, Nov. 1997.
- [51] R. Guérin, "State of the art and new trends in Quality of Service mechanisms for packet networks", this issue.
- [52] H. Schulzrinne, "IETF standards activities for telephony services: architecture and protocols (SIP, SDP, SAP, RTP)", this issue.
- [53] R. Braden, D Clark and S. Shenker, Integrated Services in the Internet Architecture: an Overview, IETF Informational RFC 1633, Network Working Group, June 1994.

- [54] S. Shenker, C. Partridge and R. Guérin, Specification of guaranteed quality of service, IETF RFC 2212, Sept. 1997.
- [55] J. Wroclawski, Specification of the Controlled-Load Network Element Service, IETF RFC 2211, Sept. 1997.
- [56] D. Hutchison, R. El-Marakhy and L. Mathy, A Critique of Modern Internet Protocols: The Issue of Support for Multimedia, in: *Proc. ECMAST '97*, Lecture Notes in Computer Science, Vol. 1242 (Springer, Berlin, 1997) pp. 507-22.
- [57] New Riders Publishing, *Advanced Windows NT Concepts*, (New Riders Publishing, 1996, ISBN 1-56205-694-8)
- [58] Data Fellows, Inc., F-Secure Virtual Private Network 1.1, White Paper, 1997. Available as <http://www.europe.datafellows.com/f-secure/vpn/overview.htm>.
- [59] W. Simpson, The Point-to-Point Protocol (PPP), STD 51, RFC 1661, July 1994.
- [60] D. Flanagan, *Java in a Nutshell, A Desktop Quick Reference for Java Programmers*, (O'Reilly & Associates, Inc., May 1996, ISBN 1-56592-183-6).
- [61] H. Kim, F. Steegmans, N. Narayanan and J. Rajahalme, Managing TINA Streams that use Internet Transport, in: *Proc. TINA'97*.
- [62] R. Mandeville and D. Shah, Gigabit Gets it Done, *Data Communications Magazine*, Feb. 1998, Available as http://www.data.com/lab_tests/ethernet.html.
- [63] Avici Systems, Inc., The World of Terabit Switch/Router Technology, March 1998, Available as http://www.avici.com/html/white_paper.html.
- [64] M. Laudach, Classical IP and ARP over ATM, IETF Standards RFC 1577, Jan. 1994.
- [65] J. Heinanen, Multiprotocol Encapsulation over ATM Adaptation Layer 5, IETF RFC 1483, July 1993.
- [66] W. Almesberger, J.-Y. Le Boudec and P. Oechslin, Application REQuested IP over ATM (AREQUIPA), IETF RFC 2170, June 1997.
- [67] D. Isenberg, Rise of the Stupid Network, June 1997, available as <http://www.computertelephony.com/ct/att.html>.
- [68] C. Smith, Applying TINA-C Service Architecture to the Internet and Intranets, *TINA'97*.
- [69] G. de Zen *et al.*, Value-Added Internet: a Pragmatic TINA-Based Path to the Internet and PSTN Integration, in: *Proc. TINA'97*.
- [70] C. A. Licciardi, R. Minerva, C. Moiso, G. Spinelli and M. Spinolo, TINA and the Internet: an Evaluation of some Scenarios, in: *Proc. TINA '97*.
- [71] F. Boujemaa (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.
- [72] J. Pavón, J. Tomás, Y. Bardout and L. Hauw, CORBA for Network and Service Management in the TINA Framework, *IEEE Comm. Mag.*, vol. 36, no. 3, March 1998, pp. 72-9.
- [73] P. Oechslin and 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://rcwww.epfl.ch>.
- [74] S. Baker, V. Cahill and P. Nixon, Bridging Boundaries: CORBA in Perspective, *IEEE Internet Computing*, Vol. 1, No. 5, Sept.-Oct. 1997, pp. 53-7.
- [75] <http://www.tinac.com/>.