Inhibitors for Ubiquitous Deployment of Services in the Next-Generation Network

Vijay K. Gurbani and Xian-He Sun, Illinois Institute of Technology
Alec Brusilovsky

ABSTRACT

The next-generation network is characterized by the merging of the Internet, telecommunications, and information technology. As envisioned some six years ago, the NGN was to revolutionize the lives of its users by providing them the best of the services of the PSTN and the Internet, written rapidly by IT programmers. Service providers were to be similarly affected by the increase in revenues paid by the users of these new services. Equipment vendors would have a difficult time keeping inventories of the NGN infrastructure. The truth has been somewhat different. In this article we characterize a set of inhibitors — technical, regulatory, and business-related — that are affecting the ubiquitous deployment of NGN services.

INTRODUCTION

A disruptive technology is one that changes an affected industry in such a way that previous competitive rules and business practices no longer apply. Internet telephony has been characterized as a disruptive technology for the telecommunications industry since it aimed to supplant the sustaining technology, the traditional circuit-switch-based telephone network.

The beginnings of Internet telephony can be traced to 1998. The Internet had by then already achieved widespread deployment. It had successfully moved from its roots in academia and commercial research labs to mainstream adoption. The two most recognizable facets of the Internet were electronic mail and the World Wide Web. With the advance of the Internet, academic research and commercial laboratories started to pay closer attention to digitizing voice and transporting it as discrete packets across the Internet.

To be sure, the idea of packetizing voice is not new. It has been researched since packet-switched networks have been in existence. What was new in 1998 were four things: first, the widespread availability of a global network in the form of the Internet ensured reachability among its participants. Second, computing power had matured to the point where it was feasible to encode and decode voice packets in real time, even in handheld devices. Third, the collective knowledge in the field of real-time transport of delay-sensitive data (like voice) was coalescing around a set of standards — Real-Time Transport Protocol, Session Description Protocol, the International Telecommunication Union — Telecommunication Standardization Sector's (ITU-T's) H.323, and Session Initiation Protocol (SIP) — that could be implemented by organizations other than telecommunications vendors. Finally, the Telecommunications Deregulation Act of 1996 created a level playing field by forcing the incumbent telephone service providers to share their equipment and network with upstarts. The combination of these four effects resulted in a shift in telecommunications. The stage was set for the disruptive technology to take over.

The rest of the article is organized as follows. The next section frames the role of services in Internet telephony. Following that, we start our examination of the inhibitors; wherever appropriate, we provide insights on how to combat the particular inhibitor, but we note that the ultimate affect of some inhibitors is hard to predict; market forces and the passage of time will allow us a clearer filter to view how some inhibitors affected the growth of the NGN.

SERVICES: THE HOLY GRAIL?

Early Internet telephony was characterized by emphasis on the media (voice in this case). Internet telephony was viewed as a means to get around paying the telecommunications operators for using their networks (a practice called toll arbitrage). If, instead, one could use one's personal computer to digitize voice and the Internet to packetize and transport it, one would not have to pay the telecommunications operators for the privilege of communicating with others. Toll arbitrage was a powerful motivator at the onset; many startups were funded to create dense port voice gateways that would convert circuit voice to packets; others were funded to demonstrate better ways of multiplexing more voice channels over transport or a better codec. However, this stage did not last for long. Incumbent telecommunication operators, sensing the threat, countered by lowering voice tariffs on local and long distance calls. This continued to the point where the rates to set up a circuit call were about the same as those for an Internet telephony call. Since the quality of the voice was much better on the circuit-switched network than on an unmanaged and best effort delivery network like the Internet, Internet telephony had to find a better answer than toll arbitrage. Thus, Internet telephony entered its next (and current) shift: emphasis on services.

Services are currently viewed as the most important ingredient in Internet telephony. We define a service as a value added functionality provided by network operators to network users. Thus, email and the Web are Internet services, while call waiting and call forwarding are public switched telephone network (PSTN) services. Early work in services for Internet telephony centered on providing users the same set of services they have been accustomed to on the PSTN [1, 2]; however, clearly, the realization was sinking in that there was a far richer set of services which could be harnessed by making the two networks work in concert.

This idea — both networks working together to provide services — quickly gained currency [3–6] and was generally viewed as an important step toward the NGN. Consider the rich services possible through such cooperation: click-to-dial (clicking on a link in a Web page would initiate a call between the Web user and a customer service representative), third party-call control (enabling a detached application to simultaneously establish a call between two or more endpoints without being part of the call), providing presence and availability for telephone lines, having the PSTN send discrete instant messages (IMs) to its subscribers on missed calls or voice mail alerts, interactive text chat between a computer user using IM and a cell phone user using short message service (SMS),or tracking the location of a cell phone user in real time and displaying it on a computer map grid.

The integration of the networks was approached from two levels: the protocol level and the application programming interface (API) level. The API proponents worked at a high level to provide the programmatic interfaces for constructing services, regardless of the actual underlying protocol that was used to communicate between the corresponding PSTN and Internet entities [3, 7]. Proponents of the protocol level worked at the detail of the actual protocol data unit that was transferred between the participating entities over the network [3, 8, 9]. In a sense, both of these approaches were complementary; regardless of an API being used, a well formatted and understood protocol data unit had to be exchanged between the entities.

Currently, it is certainly technically feasible to author and offer converged services. The technical knowhow exists as do proof-of-concept implementations [5, 8]. It is possible to perform third-party call control or to impart presence and availability of a telephone device on the wireless and wireline circuit-switched network. In addition, it is feasible to migrate an SMS into an IM and deliver it on an Internet host, as it is to participate in a click-to-dial service. It is also possible for the PSTN to send out discrete electronic mail messages or IMs to users informing them of interesting events (missed call, voice mail arrival). However, despite the demonstrated feasibility of such services, we have not witnessed a wholesale adoption of them. Outside of industry participants, the average user of the Internet or telephone network is blissfully unaware of the potential of such services.

We now take an analytical look as to why this is the case and present some inhibitors that, in our opinion, have delayed the potential of the two networks working in concert to deliver powerful services to their users.

UNCERTAINTIES

Businesses hate uncertainties, especially ones prompted by a disruptive technology. The union of the Internet, telecommunications, and information technology (IT) results in a multidomain challenge and missing expertise in all these domains. This leads to four areas of uncertainty, discussed next.

ADVENT OF THE SOFTSWITCH

The softswitch debuted in 1998 and quickly gained mindshare. The concept of a softswitch was very powerful and alluring. In its early incarnation, a softswitch was a programmable entity that insulated and acted as a signaling gateway between endpoints that could not communicate directly between themselves. For instance, an SIP endpoint would use a softswitch to set up a session with an H.323 endpoint, or with a wireline or wireless telephony endpoint. A softswitch, in this model, would also control a media gateway to set up a bearer path between the endpoints. Besides performing these functions, a softswitch was also viewed as a platform for providing personalized services such as third party call control and personalized execution of service logic applets depending on the time of day or the recipient/originator of the call. It was envisioned that services that would normally take months — or even years — to specify and deploy in the traditional telephone network would now be completed in weeks, if not days. Furthermore, since the softswitch was an Internet entity, the industry could leverage the mass of IT programmers to author services based on standardized APIs in general-purpose languages such as Java (more on this later).

However, softswitches witnessed mixed success. They were deployed, but mainly in the core of the network and not at the edges where their potential to act as a personalized service execution platform was the greatest. They did render programming of end-user service easier by hiding the details of telecommunications service complexity behind general-purpose programming languages and object-oriented frameworks. The use of softswitches as service execution platforms has forced the vendors to pay close attention to the challenges of scalability, performance, mass deployment strategies for services, and interoperability of services across different softswitch vendors. These challenges, while diverse enough, are not insurmountable. The biggest advantage of the softswitch was as a gateway to bridge mul-

Currently, it is technically feasible to offer converged services. The technical know-how exists, as do proof of concept implementations. However, outside of industry participants, the average user of the Internet or the telephone network is blissfully unaware of the potential of such services.

Internet purists who espouse the end-to-end nature of the Internet tend to view with trepidation the "walled garden" model created by service providers. This model is advocated out of necessity imposed by the need for centralized control of network resources.

tiple signaling protocols. However, that advantage was short lived as the industry coalesced around SIP as the protocol of choice for the NGN. Softswitches still exist in the NGN under the guise of application servers.

THE IMPACT OF THE CELLULAR NETWORK ON THE TRADITIONAL PSTN

Until recently, the cellular network expanded with unprecedented growth. In 1984 there were 92,000 U.S. cellular subscribers compared to approximately 140 million U.S. subscribers on December 31, 2002. The rise of the cellular network as epitomized by the Third Generation Partnership Project (3GPP; http://www.3gpp.org) was the second factor contributing to uncertainty. 3G is collaboration agreement between telecommunications standards bodies, adhered to by telecommunication service providers and equipment manufacturers.

3G blended the cellular network with the Internet like never before. It promised greater transport capacity for voice and data and worldwide compatibility. Already the cellular network was on a much shorter technology growth cycle; this coupled with the advent of voice over IP (VoIP) created an environment where business rules were greatly overshadowed by optimistic projections of the technology's advantages. Service providers spent billions of dollars purchasing wireless spectrum for future licensing, in many cases cash strapping themselves. On the wireline side, the traditional telephony vendors witnessing the rush to VoIP and 3G had to respond in kind. If the new wireless network was being constructed to be packet to the core, the wireline network had to follow suit. The traditional telephony vendors already owned the infrastructure (undersea cables, T1/E1 and optical circuits, etc.) on which the Internet ran; spurred by the optimistic projections of 3G and witnessing stagnation of their core switching business, they turned to VoIP as well. As a result, every provider had a product-specific VoIP strategy, and flush with the venture capital pouring in from Wall Street, the wireline industry (equipment vendors and service providers alike) used cheap stock to acquire the R&D, products, and customers needed to compete in a VoIP world. By 2000, the telecommunications bubble had burst, leaving cellular service providers with losses from buying expensive spectrum and the wireline service providers smarting from revenues that did not materialize as promised. Equipment vendors were also adversely affected as the service providers decreased capital expenditure.

The result of all this was the realization that converged networks and converged services were indeed possible and did provide beneficial services to subscribers. However, the uncertainty was in convincing subscribers that these services were worth paying for, and recouping investments already made in the infrastructure and booking future revenues from new networks and services. To overcome this, service providers have been looking for the "killer application" that will transform communications as we know it. Chances are that it will not be a single application that transforms communications, but rather a medley of applications working in concert to enhance the aggregate communication experience of the user.

A WALLED GARDEN OR AN OPEN OASIS?

Internet purists who espouse the end-to-end nature of the Internet tend to view with trepidation the "walled garden" model created by service providers. This model is advocated out of necessity imposed by the need for centralized control of network resources. Such a network forces signaling (and even media) to traverse a given set of centralized intermediaries. A 3G all-IP network is an example of this model, as is the current wireline PSTN and cellular network. Such service providers have complete control over the access network, and can charge subscribers for services rendered by the network. The deciding argument for centralized control is the notion that one charges only for what one controls; security and fraud prevention considerations come in a near second.

In contrast to this arrangement is the "open oasis" model of the public Internet, where all that is needed to establish a communication session is the recipient's IP address, a uniformly understood protocol, and endpoints that implement the protocol. Which view is better?

It can be argued that an incumbent NGN operator using the walled garden approach is better suited to offer existing PSTN services that use the legacy network. Thus, it can provide a seamless move toward an expanded NGN service portfolio to its subscribers by initially offering the same services as the PSTN did. Internet purists, in a need to keep the end-to-end nature of the Internet intact, reinvent PSTN services such as call transfer, conferencing, and call hold in the endpoints.

Clearly, this is an area where both views need to arrive at a happy medium. The NGN must become more than simply a PSTN replacement, and public Internet telephony must move out of recreating existing services and focus on providing innovative services for a converged network. Efforts like PINT [8], SPIRITS [9], and techniques to reuse PSTN services from Internet endpoints discussed in [1, 2] can help bridge the gap. PSTN/Internet integration — PINT — was a working group under the Internet Engineering Task Force (IETF) that proposed extensions to SIP to allow Internet entities to start telecommunications services; services in the PSTN requesting Internet services — SPIRITS — is an IETF working group that does the reverse; that is, allowing telecommunication events to be propagated to the Internet for service execution.

LACK OF INTEROPERABILITY

Interoperability between vendors has two components: core protocol interoperability and interoperability of service-related extensions. Having attended six SIP interoperability events, we note that core SIP protocol interoperability is widespread. Chances are extremely good that two biologically diverse SIP stacks will interoperate to set up, maintain, and tear down a call. Problems arise with service-related extensions. There are as many ways to transfer a call in SIP as there are to set up a third party call. Interoperability is a casualty when too many options exist to do the same thing. This is true of any application-layer protocol, including H.323.

With both SIP and H.323 becoming more complex and services becoming more important, ambiguity must be minimized in protocol extensions in order to maximize the chances that two independent implementations work in tandem.

LACK OF THIRD-PARTY PROGRAMMABILITY

The industry recognized early on that for NGN services to flourish, third party programmability was a must. It was — and indeed still is — viewed as very critical to leverage the expertise of IT programmers in the telecommunications domain.

Arguably, one reason for the success of Web services was standardized technologies such as Hypertext Transfer Protocol (HTTP) common gateway interface (CGI), HTTP servelets, and standardized APIs such as active server pages, Java server pages, and the Service Object Access Protocol (SOAP). At the onset of the Web, programming meant learning a markup language (HTML) and a scripting language(e.g., tcl or Perl) to write CGI scripts. This limited the field to a practicing few only. The Web programming industry exploded with the advent of Java. The fledgling language found a good partner in the Web. Its interpreted nature enabled Java bytecodes to be downloaded dynamically and executed in the Web browser. The language spawned later derivatives such as Javascript and Java server pages to ease the burden on Web programming. The latest addition to the Web programming field is SOAP, which is a standardized manner of encoding data between two Web entities using an Extensible Markup Language (XML) payload in an HTTP request. The instructive thing to note in the advancement of Web services is that APIs in the form of servlets and SOAP came later; the initial stages were characterized by adherents learning fairly complex and protocol-centric meta-languages like HTML and CGI.

In contrast to Web services, the telecommunication industry believed in APIs from the very beginning. The first abstract sets of APIs were defined as part of the intelligent network (IN). IN decoupled call processing from service execution by having a switch defer to a service execution platform on how to handle a call. Building blocks were specified that would allow vendors to write services that would run in conjunction with switches from different manufacturers. With the advent of the Web, the industry defined a Telecommunication Information Network Architecture (TINA). TINA further specified an open architecture (through a standard set of APIs) that attempted to focus the recent advances in the computer and networking industry on telecommunication services. As was the case with Web services, the advent of Java changed the field as the concepts of TINA were used to create new Java-based frameworks like Java API for Integrated Networks (better known by its acronym, JAIN; http://java.sun.com/products/jain/api_specs.html) and Parlay (http://www.parlay.org). These protocol-agnostic APIs attempted to leverage the mass of IT programmers to create telecommunications services regardless of the underlying network.

With the advent and use of SIP in the telecommunications domain, the focus changed from APIs to protocol-centric technologies such as SIP CGI, servlets, and an XML-based call processing language (CPL). The telecommunications industry attempted to replicate the success of Web services by defining similar constructs, including the use of SOAP as a payload in a SIP request. Other standardized protocol-centric solutions such as PINT and SPIRITS appeared as SIP-based service enablers. All of these have been successful in their own right; however, none has established itself as the solution for third-party programmability.

We expect that APIs and protocol-based approaches will coexist for a long time. We expect the standards bodies to play an increasingly bigger role here, more than they did for Web services since telecommunications services have the potential to affect more facets of daily life. Standardized solutions will lead to widespread deployment.

We expect the standards bodies to play an increasingly bigger role here, more than they did for Web services since telecommunication services have the potential to affect more facets of daily life. Standardized solutions will lead to widespread deployment.

APPLICATION OF THE WEB PARADIGM TO TELECOMMUNICATION SERVICES

In the NGN the initial stages of approaching telecommunications services mirrored those that had proved successful with the Web. For instance, building on the success of HTTP CGI, the industry defined a SIP CGI model; SIP servlets mirrored HTTP servlets; end users were empowered to create their own services through the use of CPL. Softswitch vendors even pushed the Web service model aggressively to customers including promising platforms that would download and run services in the form of Java bytecodes. However, it is instructive to note that telecommunications services are not Web services.

For one, the Web is a visual and presentation-oriented technology; its users use multimedia machines to access and enjoy the content. A Web service normally presents some information to the user for consumption. As the information is generated at the Web server, it is pushed to the browser for display. This process repeats for a finite amount of time until the user is satisfied. Telecommunications services, by contrast, do not strictly follow this model of pushing content. In contrast to the presentation nature of Web services, telecommunication services are more aural in nature. To be fair, the current generation of telecommunication devices possess a far greater ability to display images and text, but the main thrust on telecommunications services is still auditory in nature. This requires that the voice channel be present to the servicing entity such that utterances or other auditory signals (dual tone multifrequency) can be dynamically extracted from the voice channel and presented to the service logic. As it turns out, the nature of VoIP networks makes this somewhat inconvenient. Normally, VoIP signaling may traverse multiple intermediaries in order to get to its destination (think of each intermediary as a call processing

Compared to the PSTN, it is much easier to tap into an Internet network to observe, and even modify, the traffic flowing through it. Most of the traffic on the Internet is not encrypted; all that is required is a network sniffer to capture the packets and extra software to maliciously modify the payload.

switch). However, once the call is established, the media flows directly between the two endpoints, bypassing all intermediaries. Thus, if an intermediary wants access to the voice channel for services, it has to actively hairpin or trombone the call through itself. And if the service requires that the intermediary control the call (possibly to tear it down for a prepaid service), it has to trick each of the endpoints such that while the endpoint thinks it is talking to its peer, it is in reality communicating with the intermediary. Doing all this is fairly complex. There do exist solutions in VoIP that do not depend on the trombone effect, but in a public VoIP network, these would be susceptible to fraudulent billing or denial of services. By contrast, in traditional PSTN networks, the voice channel passes through every switch that handles the call, making the bearer available in a secure fashion at no extra cost to the service logic.

Yet another factor making Web services are different from their telecommunications counterpart is user expectations. A Web user encountering a site that requires the installation of a plug-in does not hesitate to download it. A few more seconds of wait time is well worth the immersive sensory experience the plug-in may provide. The same cannot be said of telecommunications services. Requiring a user to download and install a plug-in before making a call — which may be critical, a 911 call, say — is unacceptable. While in Web services the thrust is on presentation and sensory expectations, more often than not the thrust in a telecommunications service is on timing: how quickly can the call complete, or how quickly can real-time data like presence information for a telephone device be disseminated.

A final factor we consider is deployment. The deployment strategies for Web services range from individually generating and pushing content from Web servers (HTTP CGI, servlets, Java server pages, active server pages) to using a centralized repository and scheme to describe the services (SOAP and the Universal Description, Discovery, and Integration model). Even in the traditional PSTN, deploying services in a scalable and consistent manner has been a challenging aspect; while the NGN allows services to be created far more quickly, it does not aid in deploying these services to the endpoints expeditiously. In fact, a case can be made that the more powerful NGN endpoints make service deployment that much more difficult since they exacerbate the feature interaction problem and make it harder to deploy a service in a consistent manner when there are many assorted NGN endpoints, each with differing capabilities (portable personal desktop assistants doubling as phones, personal computers and laptops doubling as phones, smart cellular phones, dual-band 3G cellular and IEEE 802.11 capable phones, and finally, legacy PSTN phones and 2G/2.5G phones that still need to be supported). In a way, by allowing diverse and smarter endpoints, we increase the entropy in the system. In order to provide profitable services, network and application providers have to expend their resources to contain that entropy.

We have presented some differences between telecommunications and Web services. Despite the tendency of the industry to blur the difference between them, they are dissimilar. This is an inhibitor requiring further research to quantify the differences and suggest ways to improve both of them.

Security and Privacy

The PSTN is deemed secure partly through unintended obfuscation and partly through legal protection. Unarguably, it is much harder to tap into a telephone network and usurp signaling traffic. The communication lines and switches are physically secure from intruders. Furthermore, the inner workings of PSTN protocols, while not entirely confidential, have not been subject to as much public scrutiny as Internet protocols have been. Over the years, as society realized the benefits of a telecommunications network, legal means appeared that afforded protection to the conversing parties and made it a crime to eavesdrop without appropriate judicial authorization.

Compared to the PSTN, it is much easier to tap into an Internet network to observe, and even modify, the traffic flowing through it. Most of the traffic on the Internet is not encrypted; all that is required is a network sniffer to capture the packets and extra software to maliciously modify the headers or the payload and insert them back in the network. The situation is even more dire with newer technologies like IEEE 802.11 wireless networks. They can be compromised simply by using a wireless transceiver with close proximity (not physical proximity) to the wireless network. Clearly, security of the transport medium is paramount for users to feel confident in the network.

Privacy is another axis where the Internet falls short in comparison to the PSTN. On the PSTN, the identity of a caller can be delivered authoritatively to the called party; the service provider acts as a trusted intermediary to deliver identification information. The same cannot be said of the Internet; notice the unprecedented level of spam purportedly arriving from well-known contacts and colleagues. Privacy is extremely important in the NGN. Consider location-based services; the potential for harmful consequences exists if the location of a cellular user falls into unscrupulous hands. At best, it may violate the privacy of the cellular user; at worst, revealing the location may prove physically detrimental. Who is to blame if revealing a location escalates into a violent crime: the network provider who provided location information without consent of the cellular user? The cellular user for not setting up appropriate filters and access control lists?

Privacy in the NGN encompasses not only authentication of the peer party, but also fine-grained control over individual events that can lead to NGN services. Who should get the presence data related to a phone device? How does the system ensure that only that person receives the data and no one else? Security and privacy require a technical component (e.g., public key infrastructure), a legal framework, and cultural acceptance (having details of one's movements and call patterns made available to others). Technology can only help with the first; legal and social changes are required to overcome the other two.

OTHER INHIBITORS

Finally, we mention two other inhibitors: billing and government regulation. Billing has proven to be elusive thus far; one reason for this may be the lack of a business model. NGN operators do not yet know what the user will actually pay for. The PSTN has a well developed billing structure that should adapt itself in a manner that makes sense for NGN services. The PSTN did not envision billing for discrete services such as IM and access to presence and location events. The most basic billable entity in the PSTN is making a call. In the NGN, making a call may be a by-product of good presence indication (i.e., only if my party is present and available will I make a call). Does that mean presence will translate into premium billing? If a service provider supplies location information of a user to a vendor, who pays the service provider? The vendor? The user? The user may simply have acquiesced to allowing vendors access to her location information; she may not be willing to pay the service provider for each vendor who actually received her location. Likewise, if the vendors did not make a sale, would they be willing to reimburse the service provider?

It is instructive to note that the lack of an adequate billing mechanism has meant that the market acceptance of NGN is concentrated in enterprises and providers' long-haul networks — applications where billing is not required. The origin of the NGN is the Internet model where hosts communicate directly with each other, breaking the centralized billing model, or rendering it cost-prohibitive. NGN billing is not easy primarily due to the intelligence residing at the fringes of the network. This might be one of the main reasons why early VoIP companies resorted to "flat" monthly billing instead of a traditional metered one. Regardless of the reasons, an equitable billing model has to emerge for ubiquitous NGN service deployment.

A final inhibitor is government regulation. As of the submission of this article, the U.S. Federal Communication Commission (FCC) has decided not to intervene by regulating VoIP as traditional telephony and requiring such features as universal access, emergency services, and lawful intercept. At this early point in the trajectory of the NGN, government regulation may act as a depressor. The FCC decision has been challenged by incumbent telephone service providers, but so far the FCC has not changed its view. It is entirely conceivable that the incumbent service providers will continue to contest the FCC view. As such, any uncertainty around this decision will act as an inhibitor.

CONCLUSIONS

We have presented some inhibitors that are preventing the NGN from realizing its full potential. The good news is that VoIP has moved far beyond the hobbyist stage; commercial companies like Vonage, Net2Phone, i2 Telecom International, and open network providers like Free World Dial (http://www.pulver.com/fwd) have made VoIP a reality. The bad news is that VoIP is being viewed primarily as a technology to supplant the PSTN; however, it is far more than that. The PSTN, despite reports of its imminent demise going as far back as 1998, will be around for a number of years. An Internet-based network has the potential to act as an enabler of all types of services, in addition to voice transport, related to telecommunication. To get there, the impact of the inhibitors detailed in this article must be minimized.

Implementation and deployment of the NGN have provided us insights on new forces that may well act as facilitators to counter the inhibitors. These forces include the adoption of NGN services by enterprises, emergence of alternate transports (IEEE 802.11 and IEEE 802.16 wireless standards), the establishment of SIP as the de facto signaling protocol, and the seemingly widespread acceptance of the IMS architecture.

The PSTN will be around for a number of years. An Internet-based network has the potential to act as an enabler of all types of services, in addition to voice transport, related to telecommunication. To get there, the impact of the inhibitors detailed in this article must be minimized.

ACKNOWLEDGMENTS

The authors wish to thank the editor of the article and the anonymous reviewers for their insightful comments. Their time and effort have improved the article in numerous ways.

REFERENCES

[1] T-C. Chiang *et al.*, "IN Services for (Converged) Internet Telephony," *IEEE Commun. Mag.*, vol. 38, no.6, June 2000, pp 108–15.

[2] J. Lennox *et al.*, "Implementing Intelligent Network Services with the Session Initiation Protocol," Tech. rep. CUCS-002-99, Columbia Univ., Jan. 1999.

[3] C. Gbaguidi *et al.*, "Integration of Internet and Telecommunication Services: An Architecture for Hybrid Services," *IEEE JSAC*, Aug. 1999.

[4] C. Low, "Integrating Communication Services," *IEEE Commun. Mag.*, vol. 35, no. 6, June 1997.

[5] S. Moyer and A. Umar, "The Impact of Network Convergence on Telecommunications Software," *IEEE Commun. Mag.*, Jan 2001, pp 78–84.

[6] I. Faynberg *et al.*, *Converged Network and Services: Interworking IP and the PSTN*, Wiley, Jan. 2000.

[7] J.-L. Bakker *et al.*, "Rapid Development and Delivery of Converged Services Using APIs," *Bell Labs Tech. J.*, vol. 5, no. 3, July–Sept. 2000, pp. 12–29.

[8] S. Petrack and L. Conroy, "The PINT Service Protocol: Extensions to SIP and SDP for IP Access to Telephone Call Services," IETF RFC 2848, June 2000, http://www.ietf.org/rfc/rfc2848.txt

[9] V. Gurbani, Ed., *et al.*, "The SPIRITS (Services in the PSTN Requesting Internet Services) Protocol," IETF RFC 3910, http://www.ietf.org/rfc/rfc3910.txt, Oct. 2004.

[10] V. Gurbani and X.-H. Sun, "Terminating Telephony Services on the Internet," *IEEE/ACM Trans. Net.*, vol. 12, no. 4, Aug. 2004, pp. 571–81.

BIOGRAPHIES

VIJAY K. GURBANI [M] (vkg@acm.org) holds a Ph.D. in computer science from Illinois Institute of Technology, and M.Sc. and B.Sc. degrees, both in computer science, from Bradley University. His research interests are Internet telephony services, signaling protocols, and pervasive computing in the telecommunications domain. He holds two patents and has five applications pending. He is a member of the ACM.

XIAN-HE SUN is a professor of computer science at Illinois Institute of Technology (IIT), and is director of the Scalable Computing Software (SCS) laboratory at IIT. Before joining IIT, he was a post-doctoral researcher at Ames Laboratory, a staff scientist at ICASE, NASA Langley Research Center, an ASEE fellow at the Naval Research Laboratory, and a professor at Louisiana State University.

ALEC BRUSILOVSKY [SM] (abrusilovsky@ieee.org) has extensive experience in architecture, design, consulting, and applications development for wireline, wireless, and IP networks. In addition, he was co-chair of the IETF SPIRITS Working Group and is active within the IETF in security and Internet telephony related groups. He is is active in the IEEE Communication Society's Chicago Chapter. He holds two patents and has three applications pending.