Chapter 14

Advances in Services

James T. Smith

Consultant

14.1 Introduction

Chapter 15 Advances in Services commences with consideration of the mega-trend forces behind the explosion of new services that the communications services industry is now witnessing across all communications media—wireless, broadband—on a scale as great as, if not greater, than the recent explosion of web-based Internet-enabled services that were introduced during the late 1990's and early 2000's.

First, those mega-trends that are exerting the greatest impact on the direction that new service development is taking will be considered. Answers to such meta-questions as to how and why such mega-trend changes are expected will proposed.

The discussion begins with an extrapolation from the already observable consequences of convergence on communications architectures and infrastructures to its impact on communications services. From this consideration, a formal definition of what constitutes service convergence—of how it is characterized—will be proffered. In particular, the primal role that context management is expected to fill in the design and the deployment of converged services will be examined.

Key technologies needed for the successful development of converged communications services will be identified—generically speaking. They will be discussed from the perspective of how they contribute to the development of converged communications services. In particular, their role in the providing of context management is considered. Elements of their technical details are addressed in other chapters.

Examples of how these technologies already are being applied to transform the communications services industry will be examined in the effort to elucidate the underlying design principles that developers of 21st century communications services must master. In particular, the emerging market of small office-home office (SOHO) networks provides a rich exemplary environment for how services that exemplify the principles discussed here will be developed.

The discussion of communications services would not be complete without also considering the corresponding evolution of communications platforms that is accompanying the development of new services. The telecommunications services industry has witnessed tremendous improvements from the days of the standard black phone and later the Princess handset, and now is on the verge of a multitude of newly emerging possibilities.

The much discussed and contemplated concept of the Internet appliance is prototypical of this new communications platform development. In like fashion, this platform development is exhibiting the consequences of service convergence. In some cases, the communications platform is assuming new roles—typified by the cell phone that also has messaging capabilities, as well as a built-in calculator, day-timer, etc. On the other hand, devices not generally thought of as communications appliances are gaining such capabilities—typified by PDAs, and even wristwatches that have various forms of integrated communications support.

How well newly proposed services address these considerations will exert a profound influence in determining their potential for success. Hopefully, this higher level perspective of what constitutes successful 21st century services will prove invaluable to the student of communications service development.

14.2 Service Convergence

Several *futurist* writers have proffered their lists of century-defining mega-trends. The classic list perhaps is that provided by Donald Tapscott in his book, *The Digital Economy: Promise and Peril in the Age of Networked Intelligence*. (Tapscott 1996). Of these, the one mega-trend most apropos to communications is that of *convergence*.

The mega-trend of convergence is accomplished by the delicate balancing of the two forces of mass production and mass customization—together with a balancing of the *pull* of customer demand in juxtaposition with the *push* of technology enablement.

The element of *mass production* implies the economies of ubiquity, commonality, consolidation, standardization, interoperability, etc. In contrast, the element of *mass customization* implies dynamism, adaptability, personalization, etc. Because of the many recent technological advances, the service developer no longer has to choose one of these forces at the expense of the other. To the contrary, this counterbalance demands equal consideration be given both to service ubiquity—the manifestation of mass production, and to service personalization—the manifestation of mass customization.

The effects of this convergence take two general forms: *sustaining convergence*—leveraging the resources of mass production, and *disruptive convergence*—leveraging the creative potential of mass customization. Both of these must be considered to fully assess the impact of convergence on 21st century communications services.

These two forms of convergence—directly traceable to the general concepts of sustaining and disruptive technologies—are manifested at every level of the communications hierarchy—from the physical layer to the applications layer.

Consequences of the *convergence mega-trend* have been transforming all levels and all aspects of the communications industry—and now are beginning to impact communications services—the highest level of the communications value-chain. The focal point of this service convergence ultimately is the customer. The customer figures strategically into the core of all the sub-themes that define the *convergence* mega-trend.

To set the stage for a discussion of how convergence will profoundly influence the characterization of 21st century communications services, we begin with a formal definition of converged services, one that reflects the general characteristics of convergence just presented above.

Converged services are services that so leverage each other that the value of the resulting whole is greater than the sum of its parts. This convergence enables capabilities that the standalone component services could never deliver simply by being offered, or packaged, as independent components.

Recent examples of this service convergence phenomenon include the increased use of multi-media presentations that incorporate not only rich graphics, but also animations and sound effects, and the development of unified messaging services that attempt to manage email, voicemail, faxes, and call control as an integrated package.

Nuances of this definition will be explored in the sections that follow. First, consideration will be given to the impact of convergence on other layers of the communications hierarchy where the consequences already are more noticeable—having already affected communications infrastructure, technology adoption, etc.

14.2.1 Metcalfe's Law

The convergence mega-trend is more or less a direct consequence of *Metcalfe's Law* (Boyd 2001) which states that "the usefulness, or utility, of a network equals the square of the number of users."

Metcalfe's Law in its simplest form is a statement regarding the interconnectivity of a fully connected N-node network. There exist an order of N-squared one-to-one connections, and hence N-squared potential reasons which one could justify participation in such a network. As a network expands in geographic reach, and size—or, number of nodes, the number of connections—read opportunities existing for potentially profitable communication and interaction—grows quadratically.

Some consequences of this convergence already have been realized at the lower layers of the communications hierarchy. Technology has enabled tremendous growth of one particular network—the Internet—to such an extent that all others are being relegated to niche applications. Specialized, proprietary networks become increasingly uncompetitive and difficult to justify.

Exceptions to this pattern still exist under particular circumstances, e.g., extreme security requirements in the military establishment. Still, even such special cases will appropriate to the fullest extent possible those core technologies and capabilities of the Internet—to benefit from Metcalfe's Law to the fullest extent possible.

Such effects of convergence as these are due to the consolidation aspects of convergence. For many discussions, the concept of convergence is viewed as a process of consolidation. This consolidation includes both technologies—say, the choice of protocol standards and physical media, as well as the aggregation of businesses, and the transition or re-invention of Internet service providers (ISP) into application service providers (ASP).

14.2.2 Sustaining Service Convergence

In most such instances of convergence as those just described above, the consolidation aspects of convergence are of a *sustaining nature*—resulting in incremental improvements to core products and services, so that they operate faster, cheaper, quieter, etc. Otherwise, the fundamental assumptions remain valid—existing products and services continue to operate and contribute to the value-chain similarly to as they always have. Existing business models remain intact, and the incumbent vendors, service providers, etc. are very happy for matters to be so.

One immediate consequence of *network-level convergence* is that existing services—previously available only within particular networks—realize a degree of increased potential value purely from their ubiquity with many other potentially complementary services that could be leveraged. This consequence is compatible with a view of *sustaining technology*.

One example of *sustaining service convergence* with which everyone is familiar is typified by the universal adoption of SMTP (simple mail transfer protocol) as the preferred email protocol—over other protocols, e.g., X.400—for email interoperability. Proprietary email frameworks continue to exist—Lotus Notes and Microsoft Exchange readily come to mind, but the ubiquitous availability of email to anyone anywhere has become a *must-have* capability—an obvious manifestation of Metcalfe's Law applied to services communications layer.

14.2.3 Disruptive Service Convergence

In contrast to the sustaining convergence typified by the above described consolidation, there exists another—arguably more exciting—perspective of how convergence already has begun affecting communications products and services. This alternate perspective is due to those deeper consequences of Metcalfe's Law—a corollary:

> Until a critical mass of users is reached, a change in technology only affects the technology. But once critical mass is attained, social, political, and economic systems change.

This last observation is what Larry Downes and Chunka Mui, authors of *Unleashing the Killer App*, (Downes and Mui 1998) have termed *the Law of Disruption. The first sentence of the quote describes the world of sustaining technology*, while the latter situation is characteristic of disruptive technology.

Newly emerging communications technologies and support structures enable the user to utilize and interact with their services in previously unanticipated—and in some cases, almost unimaginable—ways. Such a serendipitous view of service creation is aligned with the characterization of *disruptive technology*.

Interestingly enough, the realm of recently emerging messaging services also provides immediately recognized examples of disruptive service convergence—instant messaging (IM) from the Internet, and short messaging service (SMS) from the wireless world. Each has become the launching point or staging area for a multitude of new value-added service capabilities that leverage the *convergence* of the immediacy of IM and SMS functionality with their core messaging functionality.

14.2.4 Finding Common Ground

Both forms of service convergence—*sustaining* and *disruptive*—depend upon the recognition of *common ground* that is shared by a variety of individual services. At the lower levels of the communications hierarchy—the physical, transport and network planes—common elements are straightforward to identify around which a case for convergence can be made.

At each layer of the communications hierarchy from the physical layer to the application or services layer, the problem space is searched for what could provide that common basis for convergence. As the focus moves up through the layers, the role of *mass customization*—customer focused, becomes increasingly pronounced relative to that of *mass production*—device and infrastructure focused.

On the one hand, the wireless industry and the PC industry will sell millions of devices, appliances, etc.—so mass production of services certainly remains an important factor. On the other hand, the multiplicity of distinct devices, applications, services, etc. with which the user may interact is rapidly ballooning.

Identification of one overarching entity in common to all of them becomes increasingly difficult. What possibly can provide the required common ground on which to establish a new framework for convergence? It is not transport, nor is it a protocol, or a software interface.

14.3 Context—Putting Communications in Perspective

The last great frontier remaining to be conquered before a ubiquitous all encompassing service convergence framework can be achieved is the identification of that entity capable of binding all services. The entity that could provide that binding, or *common ground,* for all communications services is outside the realm of communications.

The identification of that mystery entity is based on the realization that the *act of communicating* is always part of a larger *context*. The customer would be better served if each communication were adequately connected to—associated with—the context that it supports—that which gives it purpose. The synergism so derived would greatly enhance the value of such communications. Any given set of services has the potential to become a *converged service*—the whole providing more than the sum of the parts—on the basis of their support of a common context.

14.3.1 One Simple Example

With today's technology, an Internet-enabled vendor can implement an e-commerce webpage containing product offerings, along with an embedded icon—perhaps it appears as a small telephone handset—by which the customer can initiate a *separate* phone call to further discuss the contents of that webpage with a customer representative. Supposedly, this phone call could be delivered via the same PC using some form of *net-to-phone* telephony—which could be H.323-based, SIP-based, etc.

Now, what if I am upstairs on the game room PC and would like to bring my wife into the discussion? I believe she is downstairs in the kitchen, or perhaps she is now in her car traveling to pick up the kids. In any case, I do not wish to be burdened with wondering which network: home, wireline, or wireless, should be used, or wondering about the right sequence of menu clicks, keypad strokes, etc. to make the conferencing happen.

On the other hand, I want the information under discussion that provides the *context* of the phone conversation to be provided to all involved parties—to myself, to my wife, and to the customer representative. For myself, the information probably would be presented via the PC.

Perhaps, if she is in the kitchen, my wife could receive this information displayed on a wireless voice-enabled webpad, or on the flat panel display built into the face of the refrigerator. If she is mobile, then perhaps the information, in a summarized format, of course, is displayed on her web-enabled cell phone. If she is driving her automobile, perhaps she must interact verbally by means of a hands-free interface.

The primary point being made by this contrived example is that the *act of communicating* is always part of a larger *context*. Ultimately, communications services will come to exhibit what some might describe as a *self-awareness*—a sense of why they are being used, of what task is being supported, of the goal to be accomplished, to which it must contribute as a team member.

The technology to enable the integration of this level of intelligence within communications services may appear currently to be out of our reach. However, a reasonable first pass at providing such intelligence is now possible. In deed, service building blocks that could support this purpose already exist or are under development.

The customer would be better served if a communication were adequately connected to or associated with the *context* that it supports; that gives it purpose. The synergism so derived would greatly enhance the value of such communications. In fact, as later sections show, the customer already has demonstrated a profound appreciation of such context.

14.3.2 Where Is this Thing Called Context?

Most of today's telephony communications services could be characterized as *context free*, providing no real-time context regarding the purpose or the circumstance of a phone call that one is receiving. Until now, this has been the state of affairs, one tolerated by most users, with no reasonable alternatives available. However, there is more than sufficient reason to believe that the user would like to have such information, and would use it to better manage the use of the phone, or any of his communications services.

Traditionally, voice services have not been able to address the issue of context except in limited special circumstances. Perhaps this situation is due to the historical backgound of telecommunications. For decades, telecommunications services were very intelligent—human operators in the loop. The operator not only knew who you wanted, but often other contextual information critical to a call's completion. For example, the party being called may have left a note to be called at one location—say, the office. The operator, however, knew that in fact the party was at another location—say, because she had just completed another call for the party being called.

One important example where context matters is the emergency E911 service now widely deployed in the United States. The placing of an E911 call automatically implies that the call's context involves some ongoing, in progress emergency. In this case, related location information about the source of the call is presented to the agency handling it. The wireless industry currently is facing new mandates to improve the effectiveness of the contextual information provided to support E911.

Many other examples less obvious than E911 also exist. For the want of any better solution, many customers attempt to compensate for this general lack of context by employing an ad hoc combination of *caller-ID* or *calling-name-delivery* service, together with a home-based answering machine as a way to capture additional call context prior to answering a call. This really is what *call screening* is supposed to address—the gathering of sufficient context to judge appropriate treatment of an incoming call!

In many other cases, customers have been left to devise their own means to work around and to compensate for this general lack of contextual support. Alternatively, they choose to decline new services because they lacked sufficient contextual capability. In particular, the inability of currently deployed network-based voicemail services to provide sufficient real-time call context is a primary reason why many customers continue to use a home-based answering machine, even though some inbound calls may be missed, and in spite of the burden of maintaining such devices!

14.3.3 We Have a Context Logistics Problem

Context regarding a given communication—phone call, or conversation—is important. As complex as this problem may seem, it is in fact one simplification of a much greater formulation of the context problem that must be addressed: What about preserving contextual continuity from one communication to the next?

To make the general context management problem more concrete, the original example of Section 15.3.1—the webpage with the phone icon embedded in it—is again considered. Suppose that I am not satisfied with the quality or depth of the information exchanged during my first conversation after clicking the phone icon—read, *antecedent context*.

Perhaps, both the agent and I must do our homework—a little background search. For example, I may need to take a measurement, while the agent needs to seek clarification of some inadequately documented product feature. We agree to "get back to each other" with the additional information to conclude the original purpose of the webpage visit, and the supporting phone call it generated—read, *consequential context*.

In the interim—between the initial and the follow-up conversations, where is the context of the discussion *parked*? Today, that context is necessarily fragmented into separately managed pieces. This situation is an example of NO convergence! All major telephone carriers offer services by which one can park the call—place it on hold, if the hold time will not be too long. With today's telephony services, I could *park a call* if I need to pause—say to gather more information, or to await the actions of, or interaction with, someone else.

However, *any ancillary contextual information must be captured separately from the call processing.* As for capturing this additional documentation, I can make notes—write myself a post-it, or perhaps print a snapshot of the webpage.

The issue to be addressed is: how to capture enough contextual information—at any point in the process—to enable further processing of this on-going communication, without essentially being forced to retrace our steps, or even to start over. The above example demonstrates *contextual disconnect*, where the contextual chain that should connect a sequence of communications events is broken, that is disconnected, at the end of each communications session.

In contrast, with the use of *context-enabled services*, I should be able to *park the total context,* of which any phone conversations would be but a part. We should be able to resume the activity—with its total context—where we left off!

The tracking of context in such bits and pieces is the norm for today's services. But why must it be the status quo? Context-enabled services should manage my communications session as this one context, from my initial webpage viewing, to the associated phone call exchanges until finally, I am able to conclude the matter for which the communications were initiated in the first place.

For the communications industry to achieve a recognizable level of *converged services*, where the whole is decidedly more valuable than the parts, the source of the synergism that holds the parts together must be determined and addressed. That source is non-other than the *context* those parts already share in reality, and so ought to share at the service level.

14.3.4 Context Management—the Basics

The general solution to context management, as typified in the above example, is not realizable within current service infrastructures. However, much of the fundamental functionality needed to enable such services possible already is under development. Efforts that will enable the integration of such context-based features with telephony services already have been established. This section will examine those foundations.

The formal dictionary definition of **context** has two general meanings:

- *Cause-effect*—the influences and events that helped cause a particular event or situation to happen, and
- *Explanatory*—the text or speech that comes immediately before and after a particular phrase or piece of text and helps to explain its meaning.

To propose *context management* of the surrounding circumstance of a phone call, as an integral part of the phone call's *connection management*, represents an extraordinarily aggressive goal. To specify and capture all the information that possibly could impinge on a given call is beyond the scope of any system, now or ever.

However, one narrow and fairly well defined component of context with considerable utility to call management is the *setting*—timing, places, parties involved, etc. The setting should be much easier to describe and to manage than would be the total context. Currently, telephony services manage the *literal connection*—end points, circuits, packet routing, etc. They also manage simple aspects of the parties involved—hopefully, the intended parties are properly identified and connected.

The issue to be resolved is how best to proceed? How are telephony services to derive the most value from the incorporation of setting management? Perhaps the parties involved in the call's execution should be involved in the process of determining which other aspects of the setting and context are important to the call's management, and how they should also be considered in determining the best way to set-up the call? So, where are the means by which the parties can so participate?

The *causal dimension* of context—from the above definition—refers to why the call is made by the *caller* and received by the *callee*. The effect, or *consequence*, refers to the results or the benefits of the call—what was accomplished, and what could or should come next in the chain of events.

14.3.5 Context Management—an Historical Perspective

The addition of contextual support to call management is not an entirely novel, or even newly proposed, concept. This writer has published several papers on the subject, such as *"AIN System Development: The Customer Centered Service Context Profile."* (Smith 1995) As previously described, the customer has appropriated various ad hoc methods—for example, caller ID, call-blocking, and a home-based answering machine—in an effort to integrate some degree of *context management* into her *call management*.

In the mid 1990's, a number of **AIN**—*Advanced Intelligent Network*—based services were developed by the telecommunications carriers that permitted the customer to control a phone call's disposition based on such features as the date and time of the call, the identity—or at least, the caller ID—of the calling party, etc.

The new AIN service model represented significant improvements in the telecom industry's ability to deploy more quickly a greater diversity of services than previously was possible. In the theme of sustaining service convergence, much was done to create standard methods to be employed by all the telecommunications industry to provide highly portable, readily deployable advanced services.

However, the AIN approach exhibited significant shortcomings. General areas in which shortcomings can be identified include the following areas:

- The user management interface,
- Disconnectedness of service from other context management tools and services, and
- Inability to integrate the wishes of all parties—both the *caller* and the *called*.

Initially, the default user management interface was an interactive voice-response system (IVR)—whose complexity was too sophisticated yet so limited and trying for all but the most persistent users. An appropriate trade-off between providing enough flexibility, such as how many time slots in the person's schedule are enough, yet facilitating sufficient ease of use never was achieved.

Later, GUI management interfaces were developed—so that one could manage a schedule via the Web. Such GUI's were a definite improvement, but they too suffered from the disconnectedness shortcoming. Consequently, the information to be entered was often a replication of the same information already being entered into other non-telephony systems, such as the customer's company provided day-timer or planner. No one appreciated the requirement to maintain synchronization across multiple sets of systems containing the same information, especially when multiple formats, data organizations, etc. were involved.

The underlying problem was the absence of *context sharing*. As necessarily constrained by the general AIN architecture, the services were one-sided—being designed from the perspective of only one of the parties involved in the call. For example, the customer could configure his service to manage when terminating calls would be answered, accepted from whom, or redirected to where.

However, no feedback mechanism existed by which the calling party could ascertain that information—either before or during the call's set-up. No means were provided through which the involved parties could *negotiate* the terms and conditions of the call.

An often overlooked but very important value of context is its *negotiation* value. This limited state of affairs has also made all but the simplest feature interaction problems intractable. No means existed by which the involved parties could negotiate a course of action from among ambiguous or equally plausible alternatives.

14.3.6 The End Game—Converged Services Through Context

Exactly what steps and actions must be taken to achieve the quest for converged services is still an open question. Fortunately, the communications industry now recognizes the need, the challenge, the opportunity, and the urgency this quest forebodes.

John Jainschigg, Editor of the former *Computer Telephony*—recently relaunched as *Communications Convergence* in recognition of the need to respond to the *disruptive technology*-driven Internet—has described the situation thus (Jainschigg 2001):

> "Next-gen communications apps—be they vertical or horizontal in nature—require context integration before they can deliver return on investment. And context integration—adapting infrastructure and applications to drive process and express strategy—is a complicated business, requiring great technical facility, coupled with new forms of business savvy (in forms both vertical and horizontal)."
>
> "… it involves knowing how complex communications systems, workers, partners, and markets interact, both in the horizontal context of general productivity and in the relevant vertical domain. "
>
> "Today, our agents of change are building the rule-book for context integration. Because communications convergence demands this next step forward."

Note how many times the term *context* was used in the above quotation to explain communications convergence. In particular, Mr. Jainschigg in the last statement has specifically identified *context integration* as necessary to the *communications convergence* that will characterize the 21^{st} century.

So, the next major convergence milestone of the communications and information services industries is now clearly discerned. On the one hand, the communications services industry must enable context to be shared—to smoothly flow among—the multitude of applications that require it to perform their tasks.

On the other hand, the information services industry must enable the management of that context—making it available for communication as needed.

The convergence of the two—communications and information services—is much more intertwined. In particular, context serves a critical dualistic purpose—simultaneously being both the enabler of and the purpose for all communications:

- Context must be communicated as part of, or in support of a given communication's purpose—be it a phone conversation about a document being developed, an appointment being setup, etc.
- Context also is necessary to support an increasingly sophisticated communications negotiation process—from low-level call-setup, to sophisticated intelligent-agent anticipation of the customer's immediate preference—for that particular conversation.

Both the long-term goal of a context-enabled converged world, and that critical short-term immediate first step to achievement of that goal—the utilization of personal context—have been identified.

14.4 Context—the Next Killer Service Feature

Despite the shortcomings of previous efforts to integrate context management with telephony services, the concept is now more viable than ever. Which eventually will be embedded within—subsumed by—the other is still open to debate. In particular, the means to address the previously noted issues due to the lack of converged context management are now readily available.

Examples of emerging communications services that help illuminate the convergence of communications and information services around context management will be drawn from various service sectors—mobile services, service location and discovery, networked appliances, home network services, unified communications. Each example cuts across the above service sectors. Such is the nature of converged services.

Each example in the follow sections is considered from the perspective of how the service can be enhanced by and could contribute to the management of context. Context can be *momentary*: "Who's on the phone?" one might haphazardly ask. Context can be *temporal*—cumulative over time: "Call me before you leave today, so we can discuss the matter." Sometimes context is incomplete: "I guess I missed your call."

14.4.1 Presence—Personalized Context Management

One such context-rich technology now under development in the Internet is *presence management*. Jonathan Rosenberg, Chief Scientist at dynamicsoft.com, has provided the following definition:

> **Presence** is the dynamically changing set of means, willingness, capabilities, characteristics and ability for users to communicate and interact with each other.

Presence could be considered an appropriate label for that body of information that characterizes a person—it is the person's *personal context*. As such, *presence* represents a critical first step in the quest for an ultimate comprehensive treatment of context.

Interestingly, the communications industry's current interest in *presence management* is not due as much to developments in telephony, as to rapid advances in the area known as *Instant Messaging* (IM).

With the arrival of IM on the Internet, the consumer's vision of messaging has been vastly enhanced and extended. Internet-based chat programs—such as ICQ—have been around for some time. AOL and others have been very successful—able to add many new subscribers—in offering their own IM products.

In its most basic features, IM combines the functions of email—the ability to send a text message—with the immediacy of a phone call—the delivery is in real-time. This combines the two-way interactive nature of voice calls with the asynchronous nature of email, along with other feature-rich email-like capabilities—such as message archival and retrieval, and multimedia attachments.

A number of special interest groups have been organized around a shared commitment to the integration of *presence management* as a core component of communications processing. The PAM Forum[1] is one such group:

> The PAM Forum is an independent non-profit consortium with members from the voice, data and wireless networking, services and applications community. It is dedicated to establishing and promoting the presence and availability management (PAM) as an industry standard. PAM is a new, open software platform specification essential to creating advanced communications and messaging services that operate seamlessly across various telephony and Internet Protocol (IP) technologies.

The adoption of this standard by the industry will be a boon to everyone. Service providers will be able to share critical data necessary to deliver advanced, customizable services across any and all network architectures--including cable, wireline, fixed wireless, and mobile networks—including LAN's, WAN's, and PAN's. The end users will benefit from such services by being able to define, under a single identity, their personal preferences for all (participating) communication systems—including e-mail, telephone, wireless, and instant messaging.

Once we have Presence awareness, we can expect to offer location awareness for those interested in us; this leads to the obvious need to control the dissemination of information about ourselves to those without a need. Developing policy and processes to control the dissemniation of this information is a difficult problem

Working under the auspices of the IETF is the IMPP Workgroup, another effort focused on the definition of protocols and data formats necessary to build an internet-scale end-user presence awareness, notification and instant messaging system. Its initial task is to determine specific design goals and requirements for such a service. Current IMPP proposals now under consideration are: APEX (a ka IMXP), PRIM, and SIMPLE. (Disabatino 2001).

Of particular note among this group of proposals is SIMPLE. Both Microsoft and AOL have committed to the adoption of SIMPLE—essentially making it the de facto standard. The market momentum for SIMPLE—with its integration with SIP—is a reflection of the market's view that instant messaging is more than just another isolated application. SIMPLE advocates strongly envision *presence* and *instant messaging* becoming critical components of a broader suite of *integrated communications services* that includes telephone calls, voice mail and Web conferencing.

Because it is SIP-based, SIMPLE. provides a common infrastructure from which to develop a *negotiation metaphor* that incorporates the capabilities of voice communications and messaging.

14.4.2 Wireless Meets Presence Management

Additionally, the convergence of such Internet-based initiatives with the messaging capabilities of the wireless world—as typified by wireless SMS—is now a given. Ericsson, Motorola, and Nokia—major wireless equipment vendors—have collaborated on an initiative focused at the definition and the promotion of universal specifications for the exchange of messages and presence information between mobile devices, mobile services, and Internet instant messaging services.

Their Wireless Village Mobile Instant Messaging and Presence initiative will deliver architectural and protocol specifications, test specifications and tools for mobile instant messaging and presence service (IMPS). This group also will define procedures and tools for testing conformance and interoperability of IMPS. The emergence of such groups as the Wireless Village reflects through such preemptive and defensive actions the underlying concern of many service providers that Presence management represents a "sticky service" that can provide the inertia by which to hold the customers that adopt such services.

Originally a text-based service, instant messaging has evolved to include rich multimedia content, such as audio and video clips and images. ***Presence services*** provide a system for sharing personal information about the user's status, (e.g., on-line, off-line, busy), location (home, work, traveling), and the moods of their friends and colleagues (happy, angry). Presence services will allow users to subscribe to presence such as listings of which friends and colleagues are currently online. Additionally, these services will allow users to participate in private or public chat rooms with search capabilities.

Ultimately, network operators will be able to provide meeting and conferencing services with shared content. The convergence of voice and messaging services is about to attain to a new realm of functional possibilities.

The instant messaging specification will be based on prevalent bearer protocols and other well-adapted standards, such as Short Messaging Services (SMS), Multimedia Messaging Services (MMS), Wireless Application Protocol (WAP), Session Initiation Protocol (SIP), and Extensible Markup Language (XML). This service will include security capabilities for user authentication, secure message transfer and access control. Operators will find these specifications applicable to existing 2G, and the new 2.5G (e.g., GPRS), as well as emerging 3G wireless network technologies.

According to Janiece Webb,[2] Senior Vice President and General Manager of Motorola's Internet Software and Content Group,

> "Instant messaging and presence services are proving to be among the most exciting areas in today's wireless and wired world and initial signs are that this market is set to expand massively over the next few months and years."

Presence could be considered an appropriate moniker for that body of information that characterizes a person—it is the person's *personal context*. As such, *presence* represents a critical first step in the quest for an ultimate comprehensive treatment of context.

14.5 Sharing Service Context

In spite of the recent infatuation of many business strategists with Web-enabled e-commerce, etc., the most widely used and deployed Internet service has been and continues to be messaging, plain old asynchronous e-mail! Similarly, the greatest use of the telecom networks still is its "plain old telephone service" (POTS)—not fancy centrex, etc.

Perhaps an even more provocative trend to note is the relative utilitarian value that businesses now place on voice and messaging services. As early as the summer of 1998, the AMA (American Management Association), in conjunction with Ernst & Young LLP, conducted a survey to determine such preferences. Among their findings,[3] the Annual Human Resources Conference survey found that email has overtaken the telephone as the most frequently used communications tool among HR executives.

This observation is even more telling—indicative of what is to come—when one considers that the currently expanding significance of instant messaging in business was then still over the horizon.

Advances in messaging—with its generation of new requirements and capabilities, such as presence management—will be the "tail that wags the dog," rather than adding one more nuance on centrex. Messaging services provide the user with the means to share their personal context—their presence—with others as they see fit. The user is thereby enabled to take call control negotiation to levels never imagined by the telecommunications industry.

14.5.1 Messaging—an Historical Perspective

The telecom industry has been focused for some time now on just how to integrate these two lucrative services—telephony and messaging. One such approach is referred to by the catch phrase: "unified messaging," or UM.[4]

> *Unified messaging* is the integration of several different communications media, such that users will be able to retrieve and send voice, fax, and e-mail messages from a single interface, whether it be a wireline phone, wireless phone, PC, or Internet-enabled PC.

Another related effort has been focused on the interoperability of voicemail systems with each other, and with the Internet. The) initially developed a non-Internet based approach that never became widely deployed. More recently, the EMA has proposed Voice Profile for Internet Mail (VPIM), which was approved by the IETF and published as RFC 2421 in September 1998. Now, to gain the participation of a wider community, the technical work on VPIM has been transitioned to the IETF with the creation of the VPIM Work Group in early 2000.

Most of the major incumbent voicemail vendors have embraced this vision, and most carriers have attempted to introduce UM products into their markets. Web-based interfaces for their UM products have been developed. The strategists have gone on the record predicting major growth opportunities (Hicks 1999):

> "In fact, market research company Ovum Inc**.** believes that unified messaging will become almost universal in the business market, but not until 2006, with three-fourths of companies either using a service or deploying the equipment themselves. For unified messaging services alone, Ovum predicts that the number of users will increase from less than a half-million this year to 151.9 million in 2006, creating a $31 billion services market."

The future of UM—as projected by the incumbents—supposedly will be bright and shiny. The standards needed to support their UM vision are in place, and the incumbent players—the vendors and the service providers—know their parts.

But such a vision is not to be. As is the case with *disruptive technologies*, (Christensen 1999), the presumed playing field—market, technologies, major participants, etc.—already is experiencing the Internet's disruptive forces. Now the UM vision must be revisited, again.

Just as the Internet has forced the reinvention of industry segments such as electronic data interchange (EDI), instant messaging and presence management now are forcing the reinvention of the incumbent status quo vision of unified messaging, and much more!

14.5.2 Messaging—In All Its Flavors

This disruptive impact of the Internet on the telecom incumbent's vision of *messaging* is typified by comments such as the opening statement of an article appearing in *Wireless Week* (Smith 2001), "First there was voice mail, then e-mail, unified messaging, and unified communications. Now there's adaptive communications."

A plethora of messaging services have been defined in this rapidly evolving segment of communications:

- *Unified messaging* bundles voice, fax and e-mail, allowing all three to be stored and retrieved together.
- *Integrated messaging*, a subset of UM, provides the same integration but with a different interface.
- *Unified communications* provides the same bundle as UM, but gives the users real-time access. In other words, users can create a profile with preferences regarding when and where they may be reached—a user-defined *"follow me"* capability.
- *Adaptive communications* is a new variation of unified communications that learns from and adapts to the user's habits.

The above description of *unified communications*—where users can create a profile on a network with preferences on when and where they may be reached–sounds like instant messaging has arrived, and that many features previously offered as part of AIN telephony services now have been reinvented from a *communications-via-messaging* perspective.

The ADC Telecommunications announcement of *adaptive communications* introduces the concept of an *intelligent* presence agent that adapts to a person's habits, preferences, etc. based on implicit observation, in addition to the person's direct management and initial seeding of preferences via a user interface.

In its most basic implementation, IM combines the core functionality of email—the ability to send a text message—with the immediacy or spontaneity of a phone call—the interaction essentially occurring in real-time.

IM is well on its way to becoming the de facto point of convergence where the *two-way interactive immediate person-to-person* nature of voice communications is combined with the *asynchronous pause-and-start person-to-group* nature of email, along with other such feature-rich *data management capabilities* as message archival and retrieval. Furthermore, the degree of flexibility with which the user can manage and configure his use of IM already far exceeds that of the traditional telephony or voicemail services.

As previously noted, currently proposed IM standards support the use of SIP, as well as MIME, LDAP, and other Internet standards. Clearly, the environment now is ripe for the development of new classes of services that will converge the features of VoIP, IM, and more.

14.5.3 Intelligent Messaging

The *adaptive communications* proposal of ADC Communications that provides each user with her own personal *HAL*—the intelligent agent from *"2001: a Space Odyssey"*—may seem a little premature—both from a technology perspective, as well as from consideration of the customer's readiness to accept such an agent.

However, one could consider another recent announcement that leverages the IM *metaphor* and technologies.

As previously presented, the use of SIP in general, and IM in particular—as a technology—is not restricted to communications between two people. Using instant messages to interact with a computer-based intelligent agent *in the network*, rather than with another person, is an equally plausible application of this technology. In fact, new services that exploit this modified model of IM already are emerging.

One example of such a new service, called *ActiveBuddy*, is presented in the article *"Instant messages: They're not just from humans anymore."* (Courtney 2001) This service enables the user to pose a question via instant messaging to a computer—more specifically to an intelligent agent functioning as the customer's own *personal gopher* that then responds via the customer's IM client with the answer.

For example, one might request a stock quote, a sports score, or any other piece of information. In the process of determining the answer, *ActiveBuddy* might send a referral URL where more information is found, or perhaps invoke a follow-up program.

The long-term implications of this new application of IM will be quite profound. *ActiveBuddy* is the precursor of what will be a variety of new services that are enabled through instant messaging. Take a moment to grasp the significance that this expansion of IM-like functionality represents. Your PDA-enabled cell phone may come to be the *embodiment* of—or at least, an access point to—your own *personal agent:* an agent that in turn could interact in your behalf with yet other IM-enabled agents! Y*our* personal agent may as well have access to *its own ActiveBuddy* client!

As noted at the beginning of this section, e-mail gets credit for being the Internet's first *killer application*—with the Web being a strong second. Now services such as *ActiveBuddy* are atempting to combine the best aspects of both: personal communication with fast, easy information retrieval.

For the immediate future—while we wait for the ultimate realization of that grand unified convergence vision—much can be accomplished with today's resources. The necessary resources for this first convergence wave are the Presence & Instant Messaging (PIM) and Location-Based Services (LBS), together with the increasingly leveraged SIP—the *"setup"* protocol from the realm of VoIP.

In particular, the recent emergence of PIM and LBS, when coupled with the recently announced DARPA Agent Markup Language (DAML) effort to make the content of the WWW to be self-explanatory (Berners-Lee 2001)—so intelligent agents could readily process it, all provide critical resources by which to address the requirement for *context-communications integration* to facilitate true data-communications convergence.

14.6 SIP—the Oil that Makes Context Flow

The ultimate strategic value of the Session Initiation Protocol (SIP) is the ease with which it can be applied to a broad range of application domains. Certainly, SIP can be used to initiate, manage, and terminate *interactive sessions* between two or more users on the Internet. However, these *interactive sessions* need not be restricted to voice and video services, nor even to exchanges—e.g., messages—between two or more people!

More generally—and this is why SIP is such a strategic technology—SIP is an appropriate protocol by which two, or more, *endpoints*—be they people, objects, processes, intelligent agents, or some combination—to engage in a *conversation* (an interchange of *messages*) for practically any purpose. That purpose can be for most anything, including:

- voice and video,
- instant messaging,
- monitoring,
- information retrieval and exchange,
- command-n-control of applications,
- operation of devices,
- feature negotiation among intelligent entities.

As previously noted, SIP already provides the basis for SIMPLE—the proposed IM specification that has essentially already become a de facto standard, with its adoption by Microsoft and AOL. SIP also is finding its way into many other application domains—some very communications focused, and some that would not have been imagined until recently.

The wireless carriers stand to gain just as much extra value from the adoption of SIP as do the wireline carriers. The 3rd Generation Partnership Project (3GPP) is dedicated to using SIP for call control.[5] SIP is to be used from the terminal to the network, and SIP between the network service nodes. In other words, all IP Multimedia call signaling will be performed via SIP. The value of the addressing in SIP as critical to its success.

A major consequence of the decision to employ SIP as a call control method is a complete decoupling of the services provided from the network in which the user is currently operating. No longer will the limitations of the visited network have an impact on the services available to the user. Furthermore, this same decoupling effect will be realized in the SIP-enabled wireline-based network.

Another step in the direction for SIP extensions is the recently submitted IETF RFC-3087[6] *"Control of Service Context using SIP Request-URI."* This proposal would extend the application of SIP beyond the current focus on call setup and control to serve as a *bridge* between information needed to initiate and manage a communication, and additional *contextual* information that might be used by an application—perhaps, a help-desk application—that operates in conjunction, or *in sync*, with the call.

The RFC-3087 provides—as its prototypical example—the description of a voicemail service. Going beyond the previously discussed messaging services currently in fashion, this example could readily be extended to formulate a generalized approach to the convergence of voice services with messaging services. In particular, the concepts presented in RFC-3087 are suggestive that SIP could be employed as the means by which various applications could share context as some level. For example, telephony services and e-commerce services would be able to collaborate by each sharing its context as part of a converged service.

Yet another example demonstrating the potential breadth and diversity of SIP applicability is typified by another IETF draft *"Pico SIP,"*[7] proposed in February 2001. That IETF draft describes how the SIP protocol could be applied to ad hoc networks, or Pico Networks—such as one might encounter in the home environment and in mobile scenarios. This proposal leverages the direct, client-to-client, peer-to-peer session initiation capabilities of SIP.

Finally, the *meta-negotiation* capabilities of SIP are now being explored as the basis for providing a generic feature negotiation capability—as typified by the recent IETF draft *"The SIP Negotiate Method"*[8] proposed in August 2001. In this draft, the need is identified for development of a generic approach to negotiation—one that would be applicable across all protocols used in the Internet world. This has been an area of active academic research since Reid Smith introduced the *"Contract Net Protocol,"* (Smith 1980) and Randal Davis described the *negotiation metaphor*. (Davis 1983). The authors of this SIP negotiation RFC draft have attempted to attack a more immediate problem in the SIP arena and establish a generic negotiation mechanism.

14.7 Discovering Service Context

Up to this point, much has been said regarding the strategic importance of *context* to the successful development of converged communications services for the 21st century. When known, context serves as that *common ground* described in Section 15.2.4 that binds together for a *common purpose* the multiple communications capabilities of the converged service. Now, the difficult case must be addressed—namely, that in which:

- Not all the supporting or interacting components and participants of the service are known—could be devices, could be people, etc
- The context that binds them is incomplete or missing. This could include personal presence, could include details of the task that motivated the communications, etc.

Then, the missing components and contextual information must be *discovered*. Or alternately, the means and mechanisms to address incomplete situations must be determined—*negotiated*. Both of these cases are considered in the following sections.

Fortunately, much has been done to develop the *bootstrap* processes, protocols, etc. needed to effect this discovery. This is not a new concept. The classical case is the development in the United States of 411 directory services that support determining the telephone number of someone we wish to call. This once meant speaking human-to-human to request and receive the needed information. Now, this directory access has been automated with IVR front-ends that permit the calling party to speak the name of the called party and either have the number spoken, or even have it dialed directly in our behalf.

14.7.1 Discovery Characterized

As with the 411 directory service, discovery-oriented protocols, processes, etc. have been developed for each of the communications environments that have been discussed. The consideration of what needs to be discovered may be open-ended. Is it a physical device, a logical entity, some incomplete information, a location—or map of several locations, etc.? In one instance, a service may need access to a specific type of device—an Epson color printer. In other cases, any available device with printer functionality may suffice.

Likewise, the discovery process is fundamentally an open-ended process. What is determined—*discovered*—at each stage of the discovery process may well raise yet other unanswered questions. May multiple discovery paths be pursued simultaneously? The meta-search engines found on the Internet are an example of services that use this approach.

Ultimately, a service may have to proceed forward in a best effort with incomplete context. A classic example of this is found in the connectionless IP world of the Internet where there are no promises or guarantees that a packet will ultimately arrive where it was intended. For this reason, the more flexibility a service is permitted in how it proceeds, the better the likelihood that a successful completion of a given service instantiation will occur.

Often, multiple entities may be involved—each with its own parochial perspective and agenda. This situation leads to the necessity of negotiation. The negotiation that one's modem executes when one attempts to establish an Internet dialup connection is a simple example of negotiation. As the complexity, sophistication, and intelligence of services increase, negotiation will play an increasingly relevant role in the service's execution.

14.7.2 Device Discovery Potpourri

With the arrival of increasingly specialized information appliances, a number of discovery architectures addressing mobile and specialized devices have been proposed recently. These architectures are essentially coordination frameworks that propose certain ways and means of device interaction with the ultimate aim of simple, seamless and scaleable device inter-operability.

Among the better known proposals are Universal Plug and Play (UPnP), Jini, and Salutation, each of which are proposed by a major industry segment. Other minor efforts include the IETF's SLP, Lucent's Inferno, and CalTech's Infospheres Project.

Jini, developed by Sun Microsystems, provides a federation coordination framework evolved and adapted from academic research of David Gelernter (Yale University) and Nick Carriero and tailored specifically to Sun's Java.

Universal Plug and Play (UPnP), pushed primarily by Microsoft, is a framework defined to work primarily with lower layer IP network protocols, implementing standards at this level instead of at the application level, and providing a set of defined network protocols for which device manufacturers may build their own APIs that implement in whatever language or platform they choose.

Salutation, developed by device manufacturers, is drawn from research on intelligent agents and treads middle way between device autonomy and standardization., thereby enabling vendors more readily to adapt many of their legacy products to interoperate with one another.

Generally speaking, *device coordination* essentially means providing a subset of the following capabilities to a device:

- Ability to announce its presence to the network.
- Automatic discovery of devices in the neighborhood and even those located remotely.
- Ability to describe its capabilities as well as to query/understand the capabilities of other devices.
- Self configuration without administrative intervention.
- Seamless interoperability with other devices wherever meaningful.

14.7.3 Discovering and Interpreting Location

One specialized type of discovery service particularly relevant to mobile entities is *location*. For devices attached to a tethered network—typical of the office LAN, of wireline telecommunications networks, etc.—the problem of location is solved statically when the entity is installed into the network. Its static location is noted and recorded in appropriate databases—e.g., the carrier's phone directory. However, the world of communications is becoming increasingly mobile—with the cellular industry being the obvious area of growth. Mobility means location is now a dynamic entity that must constantly be re-evaluated.

The location problem is particularly critical to services such as E911 that must support emergency contexts where time to respond is of the utmost importance. Determination of where someone making an E911 call is located under all conditions and circumstances is a very difficult problem. One could argue what constitutes enough accuracy—the nearest cell tower, to the nearest intersection, to within a few meters?

Secondary concerns include such issues as privacy considerations. Should the wireless network *passively* provide the mobile party with the resources to locate himself? Then, the network only knows about the party's location when that party chooses to reveal it. Should the network actively monitor the party's location—within the provision for some agreed upon level of confidentiality?

Various technologies are now being proposed for the purpose of actually performing the location determination. Technology to locate wireless phones available today takes several forms—global positioning service (GPS) devices, overlay triangulation technologies, and cell/sector information. Other emerging approaches include ultra-wide band (UWB) technology that already is in use for radar-like search and rescue applications. The use of pico-cell technologies such as *Bluetooth* also are being considered for some applications.

Location-based services (LBS) are obviously a strategic capability to provide. In addition to emergency considerations and general tracking services, a multitude of other value-add opportunities based on LBS are being pursued. *Location-based services* are characterized by their blend of location information with other useful content—providing relevant, timely and intensively localized information directly to the consumer—on which the consumer may then act, or to yet other services that operate in behalf of the consumer.

The previously discussed presence management services seek to integrate an individual's LBS information as part of the person's presence profile. For example, a set of LBS location coordinates may be interpreted to determine that someone is at home, at work, etc., and then further applied to infer various personal states—busy, resting, available to or interested in e-commerce opportunities of a particular type—eating, playing, etc.

14.7.4 Discovery Not Limited to Physical Entities

The need for a discovery process is not limited to physical entities—i.e. devices. Sometimes information also needs to be discovered. Fortunately, methods that would facilitate such information discovery are now under development.

The World Wide Web Consortium (W3C) is studying proposed Web Services Inspection (WS-Inspection) to allow companies to post standardized directories of what kinds of Web services they have available. This specification would complement the already established Universal Description, Discovery and Integration (UDDI) standard. UDDI presents standard interfaces that companies can use to build directories of Web services that can be queried by type. UDDI performs similar to a telephone Yellow Pages: services that can be queried by business function.

Another much more ambitious effort is the Semantic Web program being persued by the W3C and DARPA, which has commenced the development of DARPA Agent Markup Language (DAML). The concept is to enable a World Wide Web that not only links documents to each other, but also *recognizes the meaning* of the information found in those documents. This is an enormous undertaking.

The first step involves the establishment of standards that allow users to add explicit descriptive tags, or *metadata*, to Web content—thereby mechanizing the ability to pinpoint exactly what is being sought. Also to be developed are methods, protocols, etc. that enable various programs to relate and to share metadata from different Web sites. From this foundation, advanced capabilities can be developed such as the ability to infer additional facts from the ones that are given. As a result, searches will be more accurate and thorough, data entry will be streamlined and the truthfulness of information will be easier to verify.

14.7.5 Negotiation Extends Discovery

A constrained negotiation framework is fundamental to each of these discovery frameworks. However, a general negotiation framework is apropos to far more than just the discovery of devices. The situation where multiple service entities are involved—each with its own parochial perspective and agenda—will almost invariably lead to the necessity of negotiation. As the complexity, sophistication, and intelligence of services increase, negotiation will play an increasingly relevant role in the service's execution.

The previously discussed proposal in Section 15.6 to develop a SIP-enabled negotiation framework is therefore quite relevant. Another recent IETF draft, *"Applying Contract net Protocol to Mobile Handover State Transfer,"*[9] from Phillip Neumiller proposes some ways the *Contract Net Protocol* (CNP) could be used to enhance macro and/or micro mobility protocols. Specifically, the draft focuses on how CNP could be used for AAA and Quality of Service (QoS) state transfer and handover negotiation in mobile IP networks.

Feature negotiation is not a new or difficult to understand capability. Similar functionality, for example, is now provided in such specific constrained domains as *smart* modems that support the auto-negotiation of multiple bandwidths, error-correction, data encryption, and data compression schemes, etc.

Negotiation is an iterative process of *announcement-bid-award* in which an entity: 1) poses its need—its requirements, or simply its preferences—to other entities that might be able to satisfy the request, 2) receives offers from responders, and 3) evaluates and accepts the best offer of assistance. Variations in the negotiation process include the possibility of iterative offers and counter-offers, of teaming arrangements by multiple responders, etc.

In each instance of negotiation, three fundamental components may be identified:

- *exchange of information* — by requests, responses, etc.,
- *evaluation of exchanged information* — by each entity from its own local perspective, and
- *final agreement* (the *contract*) — by mutual selection.

In preparation for a solution by negotiation, there are *preparatory activities* which each involved party should complete. These include such items as the clarification—from each negotiator's perspective—of:

- *absolute bounds* of this give-n-take effort,
- *compromise fall-back positions*, multiple, if possible, and
- *evaluation criteria* by which to judge any and all offers.

In support of the negotiation process, mechanisms such as languages—data structures, and protocols—enabling negotiation rules—must be in place by which each party may *accurately* and *succinctly* communicate requests, bids, etc.

Under the assumption that the involved parties are willing to cooperate, negotiation provides a reasonable vehicle for the identification of *common ground,* i.e., a *global* view of the problem's solution that is *locally* acceptable to each party involved.

In general, the system's negotiation mechanisms should facilitate the customer's indication of acceptable *fall-back* service behaviors. Much of multi-customer feature negotiation management is the matter of determination of an acceptable level of service delivery *common* to the involved parties.

14.8 The New Converged Home Network

The networked home is expected to yield a plethora of opportunities for the development of new communications services. For many, the concept of the *networked home* has referred to a home LAN of networked PCs, perhaps with attached printers, faxes, etc. Now—thanks to the new technologies being developed and price points being realized—that vision has broadened to encompass just about every appliance one could imagine—from refrigerators and microwaves in the kitchen, to the washer and dryer in the laundry, to the home infrastructure, e.g., A/C, lighting, and security. Home-based e-commerce—today restricted to a PC or possibly a cell-phone will also become embedded in every activity of the home.

Every service in the home is being re-evaluated for the additional value to be added if that service were network-enabled. For example, if your lawn sprinkler system were network enabled, then perhaps it could be managed automatically by some network-hosted weather-knowledgeable service to always water only when such is actually needed, adapting to watering moratoriums during droughts, etc. Perhaps the homeowner would like to check the contents of the refrigerator before leaving the office, or outsource the stocking the refrigerator and the pantry to a network-enabled grocery service. No doubt, yet other network-enabled services that have no current standalone, non-networked equivalent will also be created. Controlling access to services in a SOHO, or allowing interaction from without is a big issue.

14.8.1 The Current state of Affairs

Before our expectations begin to soar in anticipation of such possibilities, we first need to step back for a moment to consider the current state of affairs. The interoperability that convergence is to bring has many hurdles to overcome in the home environment. How many and what types of networks should be deployed? What management strategy should be established?

Our daily lives in our homes are touched by a multitude of distinct networks—from standalone pico-like networks to the Internet—that more or less operate as *ships passing in the night*. These include not only of the PSTN's—which are different for the United States, Europe, and throughout the world.—but also various wireless networks—AMPS, TDMA, CDMA, GSM, 802.11, Bluetooth, etc.—as well as a multitude of other explicit and implicit networks, often proprietary and non-interoperable. Traditional television—again, different in Europe and America—and radio are ubiquitous broadcast networks. The wide availability of cable-based and of global satellite-based broadcast networks, xDSL-based networks, and telephony networks is pervasive for many areas of the world. High-definition TV and digital radio are now being deployed.

A person typically may have to interact with a combination of wireline POTS phones, cordless phones, ISDN phones, any of several different types of cell-phones, fax machines, and pagers—along with various radio and infrared-enabled devices (garage door openers and keyless car doors) and appliances (TV's and VCR's).

These multitudes of often proprietary, explicit and implicit networks are non-interoperable—in terms of either the underlying communications protocols, or the information that would use those protocols. For example, information from one source—say, a speed calling list stored in my cell-phone handset—cannot readily be transferred between or synchronized with other information sources—such as with a PIM (personal information manager) on a PC or handheld PDA, or with the local telephone company's directory service.

Across the end of my own coffee table in the family room lie six different remote controls for various multimedia appliances—the TV, VCR, CD changer, DVD, amplifier-tuner, and a cable set-top box. The remote controls to my garage doors and those to my automobiles—power door locks and trunk release—are non-interoperable with each other, nor do they interoperate with my home's security system.

I keep within reach my cell-phone that I carry with me—even when at home—and a cordless phone to access my wireline POTS service—behind my key-system! I have ten distinct voicemail boxes—one on each of my family's six cell-phones, one on my work phone, and four shared by my multiple wirelines.

I have to maintain two different remote access configurations for each of my home PC's—one for access to the company's remote-access network, and one for access to the general Internet via a commercial ISP. My cable modem provides shared Internet access for a home LAN of seven PC's.

Clearly, the opportunities and the hurdles for achieving a networked home of converged services is very great. In the process, the installed base of non-interoperability such as those typified above must also be addressed. Retrofit and migration strategies will be required.

14.8.2 Great Opportunities for Service Convergence

Many developers contemplate the vision of the *networked home* broadened to encompass just about every appliance one could imagine. The question then is what are the services that this new infrastructure might enable? While a comprehensive list would clearly be a moving target with frequent additions, the following potential application scenarios are frequently mentioned.

- Home Automation and Security
- Home Entertainment
- Energy Management
- Remote Monitoring and Control
- Computer/Telephony/CE Integration
- E-Commerce

A multitude of approaches are now being proposed by various special interest groups—appliance manufacturers, entertainment providers, utility companies, etc. Until some consolidation of efforts and convergence of approaches occur, this will be quite a horse race. Of course, each of these special interest groups has formed a corresponding organization to work at defining and marketing its vision, architecture, infrastructure, standards, etc. for the networked home. Each group approaches the networking of the home from its own particular perspective—the wireless industry, the appliance industry, the multimedia industry, etc.

Some of the major organizations—alliances, consortiums, forums, etc.—announced thus far include:

- ***Bluetooth*** – http://www.bluetooth.com/index.asp
- ***Home RF*** – Home Radio Frequency Work Group http://www.homerf.org/
- ***Home PNA*** – Home Phoneline Networking Alliance http://www.homepna.org/
- ***HomePlug*** – Home Powerline Alliance http://www.homeplug.org/
- ***Home API*** – http://www.homeapi.org/
- ***ETI*** – Embed the Internet http://www.emware/eti
- ***HAVi*** – Home Audio-Video interoperability http://www.havi.org/
- ***AMIC*** – Automotive Multimedia Interface Consortium
- ***TSC*** – Telematics Suppliers Consortium http://www.telematics-suppliers.org
- ***UPnP*** – Universal Plug-n-Play http://www.upnp.org
- ***OSGi*** – Open Services Gateway Initiative http://www.osgi.org/osgi_html/osgi.html

Bluetooth and Home RF are focused on defining a wireless network infrastructure for the home. Bluetooth also addresses mobile ad hoc peer-to-peer networks—such as two cell phones swapping address book entries. Various wireless 802.11 IEEE workgroups also are considering how best to support this market opportunity.

HomePNA focuses on a network overlay of the phone wire already installed in the home, and is compatible with the ADSL-lite (splitterless-ADSL) broadband technology being delivered by the telephone companies and their resellers. In contrast, HomePlug focuses on a network overlay that utilizes the powerline running throughout the home.

Attempting to provide appropriate middleware for home network servers and appliances are Home API, Universal Plug-n-Play, and OSGi. These groups are dedicated to broadening the market for home automation by establishing an open industry specification that defines a standard set of middleware software services and application programming interfaces that enable software applications to monitor, to control, and otherwise to interact with home devices.

These groups—dominated by software and systems vendors—are dedicated to defining standard sets of software services and application programming interfaces. In particular, they allow both existing and future home network technologies such as HAVi, Home PNA, Home RF, CEBus, Lonworks, and X-10 to be more easily utilized. Furthermore, it should also be possible to integrate control of existing A/V devices (using IR-based control, for example) into one system.

In contrast, other groups with more of a hardware and component focus have offered lower-level appliance-based solutions. ETI is focused on the hardware devices that are to be made network intelligent. HAVi seeks to provide network interoperability to the multimedia devices of the home—the VCR, TV, stereo, etc. AMIC is defining standards for an embedded automobile network.

Other fruitful areas besides the home where effort is underway to embed network interoperability include in the automobile. The main players in this area are AMIC, SAE, and TSC.

The *Automotive Multimedia Interface Consortium* (**AMIC**) has announced support for the **ITS Data Bus**, an emerging hardware specification, and in collaboration with home-focused groups such as OSGi is seeking to create a common way for various electronic products to be plugged into different cars while retaining their ability to work together. For example, a navigation system, PDA, pager, and other products could share a single screen in a vehicle, with data from one item driving a response from another.

The technical foundation for a common hardware interface based on the IEEE 488 specification has been under way for some time by the *Society for Automotive Engineers* (**SAE**). In liaison with **AMIC,** the *Telematics Suppliers Consortium* (**TSC**), is developing open standards from the vehicle out to telematics services. **Telematics** is an emerging market of automotive communications technology that combines wireless, voice, and data to provide location-specific security and information services to drivers.

In addition to these consortium-led efforts, major appliance manufacturers—Seiko, Toshiba, Sony, etc.—are partnering with high-tech start-ups in their search for *lightweight* approaches to embedding Internet functionality. The focus of their efforts is to develop Internet-ready appliance components that are completely PC and OS-independent—being implemented entirely in hardware.

One such example is iReady, a startup company who is attracting much interest (Yoshida 1998) with its Internet-ready LCD's—called the *Internet tuner*. The small iReady LCD panels feature a *chip-on-flex* (**COF**) module with built-in network, e-mail and Web-browsing capabilities.

Embedding Internet-ready functionality into an appliance has the potential to facilitate applications that have little to do with accessing email, schedules, or websites. These devices, for example, could use Internet protocols not necessarily to search Web pages, but to download specific types of information available on a certain network.

The degree of service convergence that will be enabled, once a home appliance is networked within the home, is almost limitless. For example, once the refrigerator is networked, a multitude of e-commerce services are enabled—from the dynamic, instantaneous monitoring of what is in the refrigerator, to auto-replenishing from the local grocery supplier. More provocative opportunities for converged services also are enabled—as typified by medical-related services, such as assisting with assuring that someone's elderly parents have opened their refrigerator in a daily and whether they've been eating properly.

Researchers at Telcordia have proposed that the *generalized* use of SIP is in *appliance contro*l. As an ongoing follow-up to this effort, a group within the IETF is now exploring this venue for SIP at the site: IETF PAC BoF. Related drafts include *"SIP Extensions for Communicating with Networked Appliances"*[10] and *"Framework Draft for Networked Appliances Using the Session Initiation Protocol."*[11] *Figure 14.1 The SIP-Enabled Networked Home* from the Telecordia website[12] depicts how a SIP-enabled smart home would function.

SIP could support a variety of *converged sessions*. Besides call control, with SIP one could also control, program, monitor, etc. appliances. An example of a *converged service* would be to automatically turn on the lamp sitting on the nightstand beside the bed and to mute the background music—if it is on—when one answers the bedside phone in the middle of the night.

In addition to home appliances that generally are fixed in their location—the refrigerator, the dish washer, etc., another class of communications-enabled mobile appliances is emerging. One example of how the convergence of appliances and communications is proceeding is typified by the move of the watch industry to incorporate technologies to their products that allow users to communicate while on the move, find directions while hiking, access their PCs, etc.

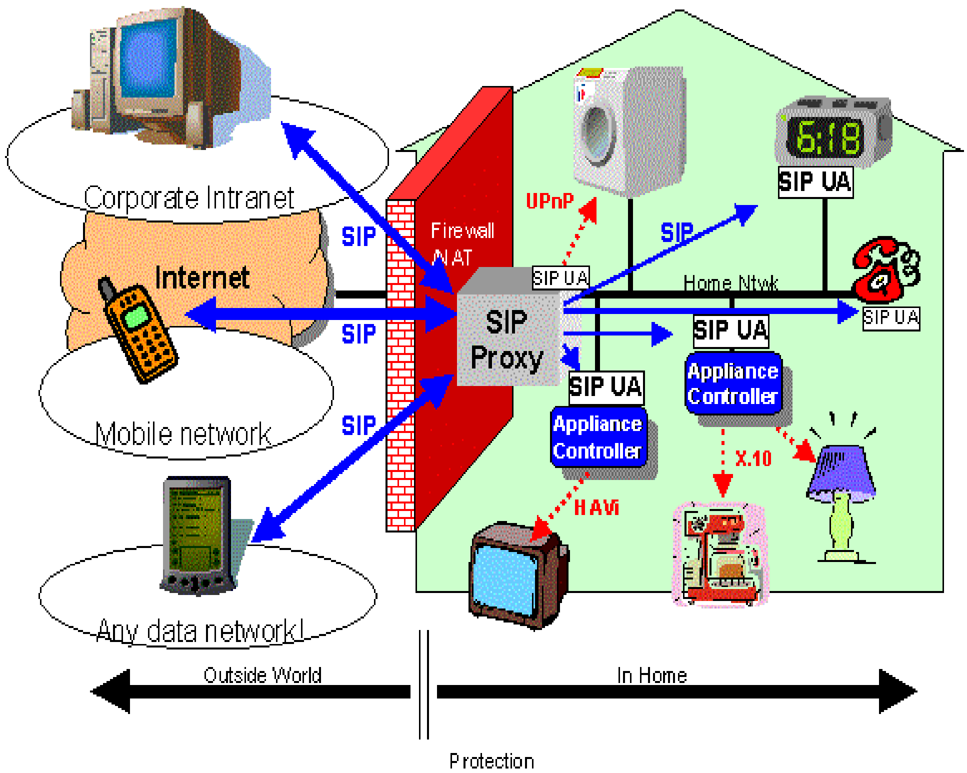

Figure 14.1 The SIP-Enabled Networked Home

First-generation versions of such devices often have employed proprietary technologies to communicate with special tethered docking stations, etc. The generation of watches now appearing in the market—e.g., Citizen's Wrist Browser wristwatch—are adopting such standards as Bluetooth—a peer-to-peer, pico-cell wireless technology. This will enable their interoperability, for example, with a variety of other communications-enabled devices—from the new generation of Bluetooth enabled cell phones, to Bluetooth enabled cash registers and vending machines, as well as Bluetooth enabled VCRs.

How many of the population will opt for Dick Tracey wrist phones, versus how many will choose to wear cell phone-friendly wrist watches, or yet some other approach to service and device convergence, only time will tell.

14.9 Conclusion

This chapter has attempted to provide the reader with a panoramic view of the almost limitless possibilities that the convergence mega-trend is now bringing to the realm of communications services. The *out-of-the-box* thinking that will be required of developers who would take full advantage of the emerging technologies that are enabling this convergence has been described and hopefully meaningful examples from which the developer could build and extend.

The great news is that the realization of not only these but many other possibilities that enable and leverage converged communications services are now upon us. We are living in exciting times!

References

Berners-Lee Tim, Hendler Jim, and Lassila Ora, 2001 *"The Semantic Web,"* Scientific American, May 2001.

Boyd Charles 2001 Why Strategy Must Change, *Management Issues: MGT 487*, SMU.

Christensen Clayton M. 1997 *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business School Press.

Courtney David 2001 Instant messages: They're not just from humans anymore, ZDNet, May 1, 2001.

Davis R and Smith R G. 1983 Negotiation as a metaphor for distributed problem solving. Artificial Intelligence 20,1,63-109, 1983.

Disabatino Jennifer (2001) IM group works toward interoperability, Computer World, February 8, 2001.

Downes Larry and Mui Chunka 1998 *Unleashing the Killer App—Digital Strategies for Market Dominance*, Harvard Business School Press.

Hicks Matt 1999 Messaging systems of the world, unite, PC Week, October 4, 1999.

Jainschigg John 2001 Agents of Change, *Computer Telephony,* April 2001.

Rosenberg Jonathan 2001 Evolution of IP-Based Presence Services, dynamicsoft.com.

Smith Brad 2001 Yes, Mom, I'll Take Your Call, Wireless Week, April 2001.

Smith James T 1995 *AIN System Development: The Customer Centered Service Context Profile*, Telecom Information Networking Architecture, TINA-95.

Smith R. 1980 The contract net protocol: High-level communication and control in a distributed problem solver. IEEE Transactions on Computers, C-29(12):1104-1113.

Tapscott Donald 1996 *The Digital Economy: Promise and Peril in the Age of Networked Intelligence*, McGraw-Hill.

Yoshida Junko and Hara Yoshiko, 1998 *"iReady, Seiko develop Internet-ready LCD's,"* EE Times, October 26, 1998.

Index

Glossary and Acronyms

3GPP	3rd Generation Partnership Project
AIN	Advanced Intelligent Network
AMIC	Automotive Multimedia Interface Consortium
AMPS	Advanced Mobile Phone System
APEX	Application Exchange
AS	Application ServiceProvider
CNP	Contract Net Protocol
CDMA	Code Division Multiple Access
DAML	DARPA Agent Markup Language
DARPA	Defense Advanced Research Projects Agency
E91	Enhanced 911 (emergency service)
EMA	Electronic Messaging Association
ETI	Embed The Internet
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communication
GU	Graphical User Interface
HAVi	Home Audio-Video Interoperability
Home API	Home Application Programmer Interface
HomePlug	Home Powerline Alliance
Home PNA	Home Phoneline Networking Alliance
Home RF	Home Radio Frequency Work Group
IET	Internet Engineering Task Force
IM	Instant Messaging
IMP	Instant Messaging and Presence Protocol
IMPS	Instant Messaging and Presence Service
IMX	Instant Messaging Exchange Protocol

ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
IVR	Interactive Voice Response
LAN	Local Area Network
LDA	Lightweight Directory Access Protocol
LBS	Location-Based Service
MIME	Multipurpose Internet Mail Extensions
MMS	Multimedia Messaging Services
OSGi	Open Services Gateway Initiative
PAM Forum	Presence and Availability Management Forum
PAN	Personal Area Network
PC	Personal Computer
PIM	Personal Information Manager
PIM	Presence and Instant Messaging
POTS	Plain Old Telephone Service (or System)
PRIM	Presence and Instant Messaging
QoS	Quality of Service
SAE	Society for Automotive Engineers
SIMPLE	SIP for Instant Messaging and Presence Leveraging Extensions
SIP	Session Initiation Protocol
SLP	Service Location Protocol
SMS	Short Messaging Service
TDMA	Time Division (or Demand) Multiple Access
TSC	Telematics Suppliers Consortium
UDDI	Universal Description, Discovery and Integration
UM	Unified Messaging
UPnP	Universal Plug and Play
UWB	Ultra-Wide Band
VoIP	Voice Over IP
VPIM	Voice Profile for Internet Mail

W3C	World Wide Web Consortium
WAN	Wide Area Network
WAP	Wireless Application Protocol
WVMIMP	Wireless Village Mobile Instant Messaging and Presence
xDSL	Digital Subscriber Line (of any type)
XML	Extensible Markup Language

Web Sites

1 The PAM Forum, (2001) http://www.pamforum.org/

[2] Mobile Messaging Standards Coming, AFX News Ltd., April 26, 2001, http://www.mbizcentral.com/story/m-allNews/MBZ20010426S0004

[3] E-MAIL BEATS TELEPHONE AS TOP MEANS OF WORKPLACE COMMUNICATION ACCORDING TO Ernst & Young/HRFOCUS SURVEY, (1998) http://markbushstudio.com/hk/ernieweb2/About/ey10.html

[4] Unified Messaging Tutorial, International Messaging Consortium, 2001 http://www.iec.org/tutorials/unified_mess/

[5] Third Generation Partnership Project, 2001, http://www.3gpp.org/

[6] Control of Service Context using SIP, IETF RFC-3087, April 2001, ftp://ftp.rfc-editor.org/in-notes/rfc3087.txt

[7] Pico SIP, February 2001 http://search.ietf.org/internet-drafts/draft-odoherty-pico-sip-00.txt

[8] The SIP Negotiate Method, March 2002 http://www.ietf.org/internet-drafts/draft-spbs-sip-negotiate-01.txt

[9] "Applying Contract net Protocol to Mobile Handover State Transfer," Phillip Neumiller, Nov. 27, 2000. http://www.ietf.org/internet-drafts/draft-neumiller-seamoby-cnpmobility-00.txt

[10] SIP Extensions for Communicating with Networked Appliances, November, 2000 http://www.argreenhouse.com/iapp/draft-tsang-sip-appliances-do-00.txt

[11] Framework Draft for Networked Appliances Using the Session Initiation Protocol, November, 2000. http://www.argreenhouse.com/iapp/draft-moyer-sip-appliances-framework-01.pdf

[12] AR Greenhouse, 2001 http://www.argreenhouse.com/iapp/