

Service-Oriented Research Opportunities in the World of Appliances

Vijay Machiraju, Martin Griss, Klaus Wurster,
Reed Letsinger, Akhil Sahai, Aad van Moorsel

HP Labs
Software Technology Lab
Palo Alto, California, USA
aad@hpl.hp.com

1 Introduction

There is a rapid growth in the number of ways that appliances are being put to use. New and novel services are being created, personalized, and then delivered by exploiting e-services and appliances. Appliances are being used not only as service clients but also as interfaces to personal agents that help users make better decisions. On the other end, technologies such as Bluetooth are enabling appliances to communicate over adhoc networks and to act as service providers for a new class of services.

Building, deploying, and using such services requires interplay between a number of technologies – service composition, registration, brokers, discovery, introspection, binding, agents, wired and wireless protocols, personalization, and context-awareness to name a few. Emerging high-level middleware platforms such as the XML stack in HP Netaction/Bluestone [1], Biztalk in .NET [2], and place managers in CoolTown [3] are aimed at providing some of these features.

While the addition of appliances brings a whole new class of services into place, the problem of managing these appliances and end-to-end services becomes more challenging. For example, service providers of these new services will be faced with the problem of having to manage a very heterogeneous environment consisting of multiple networks, protocols, and devices. Users of appliances will need additional utilities such as a profile management, global login, and security.

This document is a collection of thoughts, interesting technical problems, and opportunities for service-oriented research in the world of appliances. Services in this world could mean (Figure 1):

- (a) Network and Infrastructure services that would enable the creation, deployment, operation, management, optimization, and evolution of business services. These are horizontal services that would cut across and would be used in multiple domains.
- (b) Vertical business services (domain-specific) that are built for specific markets.

This document focuses on two very broad areas of research:

1. What are the new types of business services that are enabled by appliances, or by a combination of appliances and other existing technologies such as e-services?
2. What are the network and infrastructure services that would enable the creation, deployment, operation, management, optimization, and evolution of such business services?

The rest of this paper is organized as follows: Sections 2 and 3 provide some groundwork to get to a common understanding on appliances, their trends, usage scenarios, and characteristics. Sections 4 and 5 discuss opportunities for research in the two areas listed above.

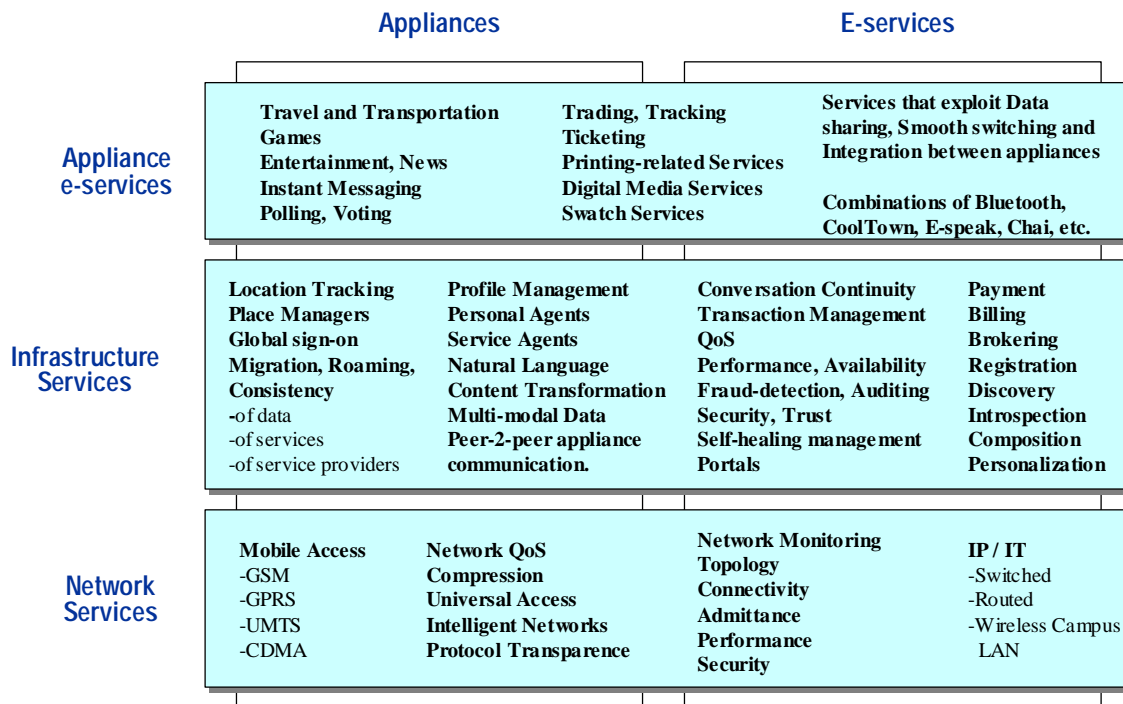


Figure 1. Appliances and Services

2 Appliance Scenarios

This section provides some scenarios that demonstrate how appliances are being used or are anticipated to be used. Understanding these scenarios would help in the rest of the discussion.

1. **Personal utility:** Appliances such as the earlier versions of Palm® and notebook computers have personal utilities such as address and phone book, notepad, etc for assisting users in managing their day to day activities.
2. **Thin service client:** This is a popular way that appliances such as cell-phones and Palm handhelds are being used today. They act as simple browsers that connect to special services on the Internet. For example, many cellular service providers offer (or contract other service providers to offer) stock trading or banking services to their customers. In all these cases, the appliance displays a simple browser (a general purpose software) or a special-purpose user interface to provide access to the service's functionality. In other cases (e.g., games), the appliance may download special-purpose software from the Internet.
3. **Interface to personal agent:** It is anticipated that the world in the future will comprise of a sea of agents – some of which collaborate to provide services (e-commerce agents), and others which act on behalf of end users (personal agents). These agents communicate with other agents and with existing resources on the Internet to accomplish their tasks (Figure 2). For example a personal agent may help the user in selecting the best service by discovering and negotiating with a broker agent. Similarly, a personal agent may interact with the shopping mall agents and with the store agents (in a shopping scenario) to provide the best options and deals to the end user. In all these cases, where pervasive access to the personal agent is required, an appliance acts as the interface to the personal agent. This scenario is not very different from the “thin service client” scenario above from the appliance's perspective. But, it stresses the new class of services that could be offered when appliances are combined with agents and services.

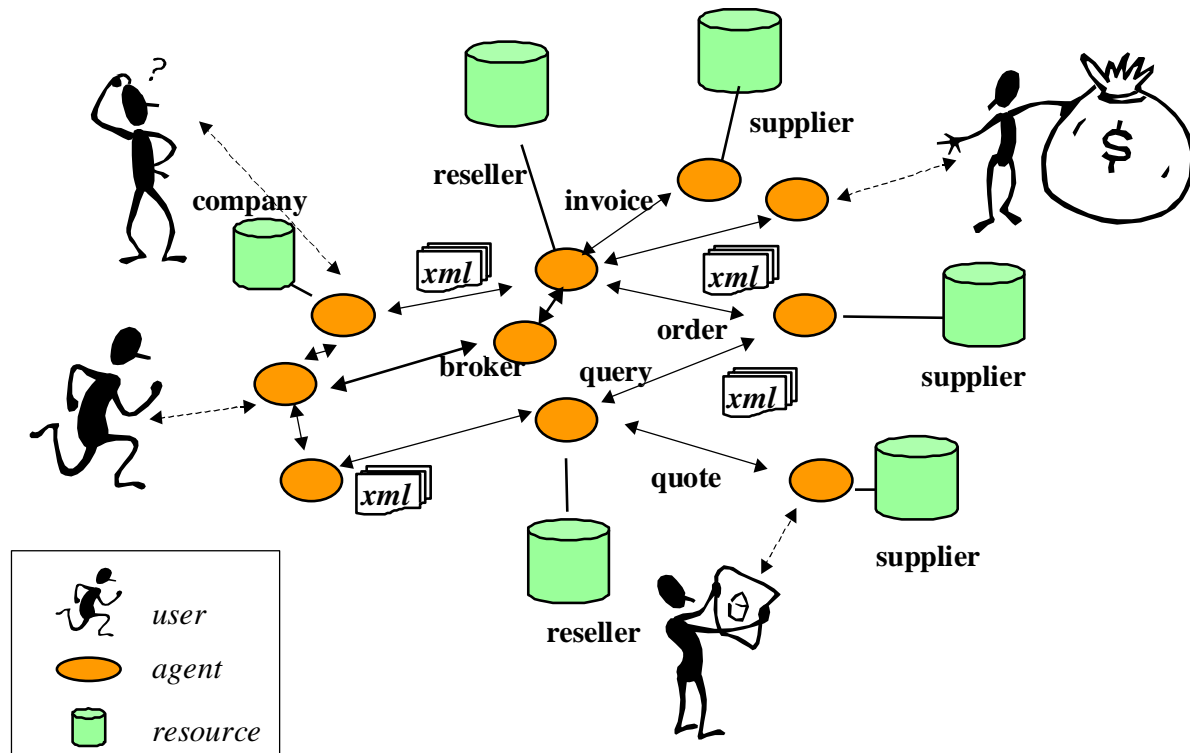


Figure 2. E-commerce agents, personal agents, and users interact with each other.

- 4. Embedded web server, identifiable:** Some of the appliances (including those that are fixed) may host embedded software such as a web server. In other cases, these appliances may just provide a unique URL (either using a barcode mechanism or by squirting) to another appliance, which can then be used to connect to a web server located elsewhere. In the simplest case, this enables web pages to be served by the appliance. These web pages could provide basic information about the appliance or even expose some configuration interfaces. In the more advanced case, the web server could host complex objects or agents (e.g., Chai Web Server), thus allowing appliances to download new features or to support complex conversational interfaces.
- 5. Network of appliances:** Appliances can communicate with each other. For example, a handheld can beam a document to be printed to a printer. An appliance may use a cellular link to download the profile of its user. In all these cases, appliances exploit the network to provide a new service to their users. In other cases, appliances communicating on a network may offer a new class of services to other customers thereby bringing new sources of revenue for the appliance users. For example, appliances on a adhoc network can exploit each others' unused resources (memory, CPU, connections) etc to perform a better job at what they do. An appliance downloading a huge movie on a cellular link could take a long time to finish the job. But, if the same appliance could exploit the cellular connections of the proximate appliances, it could easily speed up the download process. The fast adhoc network could then be exploited to bring all the pieces of the downloaded movie into one appliance. In examples such as these, there should be payment mechanisms between the borrowers of the appliance resources and the providers of such resources.

3 Characteristics of Appliances

- **Heterogeneity:** Appliances have widely varied capabilities and configurations. Differences exist in terms of appliances' form factors, functionality, mobility, user interface capabilities, connectivity, and protocols used. Further, a device's capabilities are subject to frequent change, yielding frequent change in interfaces and communication protocols.
- **Special purpose vs. generic:** Most of the appliances are made to be special purpose (e.g., toasters, cell phones, etc). Some of them are intended to be general-purpose clients (e.g., handheld browsers).
- **Large number:** A feature that makes appliances so different is that there could be so many of them in use. If we consider the number of special purpose appliances and those that are mobile, we can imagine how many of these there could be.
- **Mobility:** Some of the appliances are mobile with the user (e.g., watches, handhelds, etc.). Some of them may be mobile in other ways (e.g., with the user's car).
- **Intermittent connectivity:** Some of the appliances (e.g., PC) are always connected. Others are intermittently connected. For example, a handheld may be connected when it is plugged in. The rest of the time, it is disconnected. Some of the appliances may have multiple connections. For instance, an appliance may have a cellular link, a Bluetooth link, and a traditional IP link. These connections differ widely in their characteristics. Cellular links are long-range, low bandwidth, and expensive. Bluetooth links are short-range, high bandwidth, and cheap. IP links are medium-range, high bandwidth, and cheap as well. Cellular links and IP links are more pervasive than the Bluetooth links.
- **Adhoc networks:** The networks that appliances participate in may be adhoc in nature. Networks are dynamically created and broken based on the proximity and locations of devices. The only people who can see this network are the users of the appliances. It should be possible to manage these networks and applications built on the top of these networks without the need for additional operators, or even better, make them self-manageable.
- **Context dependence:** The functionality offered by an appliance may change based on the context in which it is used. For example, based on who is using the appliance (user-dependence) the appliance may load different profiles. Based on the location in which it is used (location-dependence), the appliance may access new services.

The services that are built around appliances also exhibit unique characteristics:

- **Heterogeneous networks:** In the end-to-end case, the service may have to pass through various network infrastructures such as IP-based, telco-based, wireless, and adhoc networks.
- **Heterogeneous middleware:** Components of the service will be implemented on various middleware platforms such as Bluestone and .NET for composed e-services and agents, CoolTown for bringing context awareness, and appliances for clients.
- **Intermittently connected providers and consumers:** Because of the intermittent nature of connectivity of the appliances, providers and consumers of services may not always be connected to each other.
- **Context-dependent access, functionality, QoS, and performance:** Based on the context (location, user, etc), a whole new set of services may be offered. Non-functional characteristics of the service such as QoS and performance may also vary widely based on the context.
- **Security and Access:** Providing secure communication and access to services will be even more challenging with the large number of appliances, some of which could be mobile. There are various issues surrounding authentication, authorization, access control, and trust when appliances are used to connect to services.

- **Pervasiveness:** Appliances facilitate a pervasive computing infrastructure where consumers can have access to information and services from any location. This puts a tremendous burden on how such services could be built and how scalable, available, and manageable they should be.

4 Opportunities in New Business Services

Figure 1 showed a number of business services that could be created by using appliances and other e-service infrastructures. Some of these are explained below using a characteristic scenario of how such a service could be used.

- **Printing-related Services:** Carrying documents or references to documents on appliances and printing them on a proximate printer in airports and other locations.
- **Digital-media Services:** Pictures in a camera being uploaded to a nearby printer for printing.
- **Games:** Service providers offering multi-user games on appliances.
- **Shopping:** A combination of personal agents (representing the user) and service agents (from the mall and shops) presenting a unique shopping experience to the user through his or her appliance.
- **Smooth move-over between appliances**
- **Appliance striping:** Users using unused resources (CPU, memory, connections) on each others' appliances.

5 Opportunities in Infrastructure Services

These are services that could be considered to be part of the infrastructure. Their role is to facilitate the development, deployment, operation, and management of the business services. For the new world of services, it is important to stress that the dynamic "run-time" or "provisioning time" binding and customization of components will change the nature and role of management, and the pace of evolution of components and composed systems. Traditional management aspects such as monitoring and control will remain in this new world. At the same time, management could take a different form - in the form of services - which will make it to be an integrated part of the infrastructure.

5.1 Development, Deployment, Operation

- **E-commerce agents and personal agents:** Agents will play a big role in enabling the next generation of services. A set of agents called e-commerce agents will be good at B-2-B conversations and will help in the creation of composite e-services. Another set of agents called personal agents will help end users in C-2-B conversations. Users will interact with personal agents through appliances [Figure 3].
- **Messaging:** It is hard to imagine a world where all the agents and appliances will have a static set of features. One way to make this more dynamic is to add the capability to learn (download) new features on demand. One way to accomplish this is to add general-purpose software such as a web server to each agent or appliance that is capable of receiving a message (in XML, say). Based on the type of message that it receives, the agent or appliance can download new software/features that can be used to interpret those messages (e.g., the XML message includes information on the DTD and its interpreter). This is similar, but an extension, to the current mechanism of extending MIME types supported by a web server/browser by adding new plug-ins.

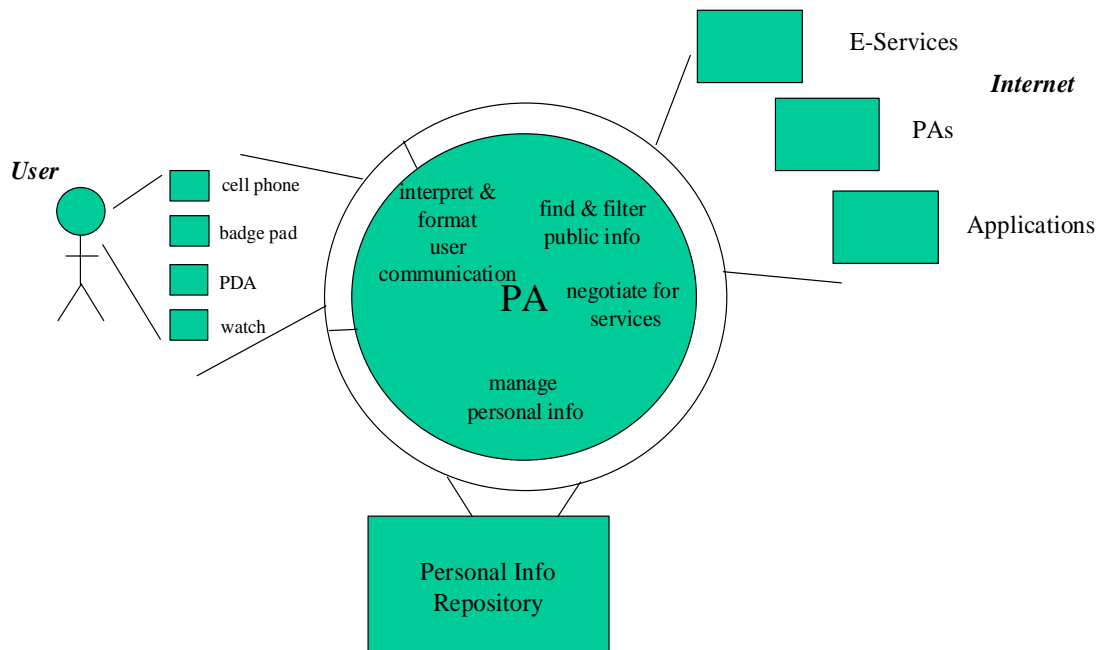


Figure 3. Appliances interact with personal agents, which then communicate with e-services and other applications.

- **Publish-subscribe communication:** To facilitate scalable, decoupled communication that can support advertisement, discovery, and usage of services, we need a flexible infrastructure such as publish-subscribe/multicast/federation of buses. This will take us a long way in dealing with the large number of appliances, the numerous types of messages, and the dynamic nature of these services.
- **Netaction + CoolTown:** While Netaction is good at facilitating the creation and composition of services on the Internet, CoolTown is good at adding context-sensitiveness to the information and to the service. A combination of these two technologies can deliver powerful context-sensitive services and appliances can be used as the means of delivering such services [Figure 4]. For example, a CoolTown beacon located at a shop can squirt a URL to a Bluestone service, which can then provide information about various products or even an automatic purchasing service. The appliance could provide the credit card information to the e-service which then uses a payment e-service for authentication.

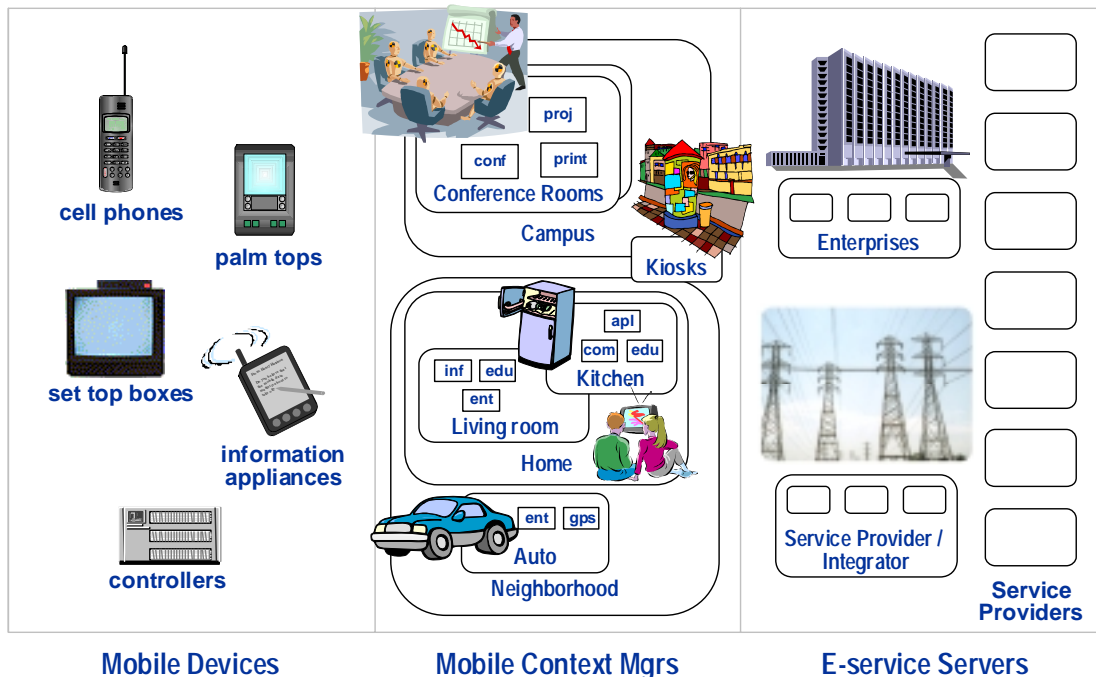


Figure 4. E-services and context managers can be combined to create a new generation of personalized/context-sensitive services that can then be delivered through appliances.

- **Bridge gaps between Netaction, CoolTown, wireless, .NET, Bluetooth...:** These technologies cover ground in various aspects in an end-to-end scenario. Some of these provide middleware infrastructures (Bluestone in Netaction). Some of these define protocols for communication or document exchange (e.g., XML stack in Netaction, SOAP for .NET). New standards such as UDDI, eCo, CBL, emerge every day. Mechanisms for bridging the gaps between these protocols should be designed without having to reinvent yet another protocol.
- **Natural Language**

5.2 Management

Listed below are a few challenges that arise in terms of managing appliances and services that appliances participate in. Management solutions that address some of these challenges could be targeted for the users of appliances and/or service providers.

- **Integrated management of the heterogeneous environment:** One observation from the OpenView efforts is that Telco style management and services and application/service style management are being integrated. Providing a uniform interface to the service provider or to the user while hiding the heterogeneity in the appliances, networks, and infrastructures is quite challenging.
- **Management-aware appliances and services:** One idea that has emerged is that traditional OVBU style application and service management changes or is extended to much more "self-managing" or "management-ready" components. These components will incorporate or expose management interfaces and capabilities, rather than being interfaced to a monolithic management system. This built in management can be seen as a service (or service facet) that can be dynamically composed with the service facets of other components, and can be used/presented through a management service and portal.

- **Tighter integration between management and the service:** Management services can now become a source of immediate information to developers / maintainers / evolvers. They can be used sort of like profilers for compilers or as optimization and brokering utilities to pinpoint and quickly change system hot spots, or dynamically tune components or the composition.
- **Location transparency and visibility:** There will be new opportunities for management capabilities that help in location related services; some capabilities will enhance and support location-aware modalities, while others will work together to conceal location and mobility. For example, as a user moves around, the management service component will cache, replicate and relocate some elements of the system to maintain QoS, hide transitions across multiple service providers, adjust and negotiate tiers of service as different agents and services appear in the same environment. Furthermore, the personal agent (or something use the profile) will maintain negotiations, and participate in task-prioritized QoS adjustments as the mobile user moves, or periodically connects and disconnects.

Likewise, location-aware capability will find and acquire different services as the mobile user (or the users agent) move to new locations, and will need to manage those, the agents and possibly the appliances and available resources. Some of these capabilities are already available in Telco/mobile environments at a lower level and even some user level services (e.g., rule/role/mobility based forwarding of calls between phones, mobile phones and pagers), but can be extended into the higher-level services that can become location aware.

Some of these are explained in more detail below:

- **Moving data/service along with the user (Follow-me):** As the user moves around, the server closest to the user should serve the user. There are host of issues involved in user management, user data migration etc.
- **Content transformation:** Depending on the mobile device resources, the data has to be transformed to suit the bandwidth, presentation capabilities, and resources available to the user as he/she moves around.
- **Location transparent QoS and performance:** The service provider may want to guarantee quality of service and performance parameters as a function of various parameters such as location, context, and connectivity.
- **Visibility and control on data/service migration:** The end user may want visibility and control over each of the above characteristics. For example, a user may prefer one server to another when presented with a choice.
- **Management of user data:** The data the mobile user has, soon becomes stale. There are host of issues involved in data consistency. The changes made by the user have to be reconciled with the actual data. There has been lot of research in this area (Dr. Satyanarayana's work on cache coherency in mobility).
- **Profile Management:** Users need to be able to create and edit their personal profiles. These profiles consist of user's preferences and policies. Users also need to be able to control/restrict visibility of information from their profiles to other service providers or agents.
- **Mobile managers:** The managers managing appliances and services could themselves be mobile.
- **E-2-E availability and performance:** The service provider has to guarantee end-to-end availability and performance of the service despite the intermittent connections. For example, a user may move to a location expecting to be connected to the service. If the user cannot establish the connection when he is expected to, the end-to-end service is affected.
- **Management protocols for dealing with intermittent connections:** The device may not be connected when it faces a problem. The management system should be able to collect information for diagnosis from the appliance when it is eventually connected. Similarly, the management system may not be able to fix the problem with the appliance immediately. It may have to wait for the

appliance to be connected. Another way to deal with intermittent connectivity would be to have a distribution of agents for managing the appliance. Some of these agents could be located in places where the appliance could come into contact with them. The agents can then collaborate to fix the problem.

- **Transaction management:** The user may disconnect willingly or may get abruptly disconnected because of lack of bandwidth or battery. There are host of issues regarding rollback, specially so for transactions the user was about to commit or the conversations the user was having with a service (say, while using voiceXML).
- **Selection of appropriate channels:** To optimize the costs and performance, the user could have some control (e.g., policy-driven) on which communication channel is most appropriate for a particular task.
- **Operatorless-management:** In the case where appliances enter and break out of adhoc networks, there is no operator for performing network or service management tasks. The end users are themselves responsible for tasks such as handling failures, admittance policies, and topology/connectivity management. In such cases, self-healing, adapting appliances and services will be of great benefit.
- **Global sign-on:** With all the vast number of appliances and mobility, users would need the facility of global signon. Its authentication information has to be verified as and where the user logs in.
- **Software updates:** Appliances should be able to download new versions of software as and when them become available or are needed.
- **Managing personal agents:** When appliances act as interfaces to personal agents, users should be able to control the behavior of their personal agents through appliances.

6 References

[1] Netaction and Bluestone. <http://www.hp.com/products1/softwareproducts/software/netaction/index.html>.

[2] .NET. <http://www.microsoft.com/net/default.asp>.

[3] CoolTown. <http://cooltown.hp.com>.