Moving from Internet Appliances to Internet Intelligent Environments: Challenges and Directions

Edwin O Heierman, III, Diane J. Cook, Sajal K. Das, and Lawrence B. Holder
Department of Computer Science and Engineering
The University of Texas at Arlington
Arlington, Texas 76019-0015
{heierman, cook, das, holder}@cse.uta.edu

ABSTRACT

Although various research efforts have reported some success in solving difficult problems related to intelligent environments, very little progress has been made in implementing a truly intelligent system in reality. This paper looks at intelligent environments from a networking and connectivity point of view in order to gain a fresh perspective. First, we review intelligent environment qualities and projects so that we may ascertain the current state of intelligent environment research. Next, we present a classification of the essential components of an intelligent environment, and observe that a connectivity infrastructure does exist in the form of a network of Internet Appliances. We also note, however, that the intelligence components necessary for an intelligent environment are lacking. Therefore, we discuss challenges and provide directions for future research so that we can move from a network of Internet Appliances to Internet Intelligent Environments.

1 INTRODUCTION

Since the beginning of the computer age, extraordinary predictions have been made about computers revolutionizing our lives. Early predictions were based on the perception that computers could eventually perform the same tasks as humans, and possibly even exceed human intelligence. Although computer performance has significantly increased over the years, many of these predications have failed to materialize due to the difficulty of learning even simplistic human tasks. For example, consider the prediction of intelligent home automation. Widespread automated homes do not exist to date. For most of us, perhaps a VCR that could program itself would be a significant step towards an intelligent home.

Two observations clarify the difficulties with intelligent environment systems. First, an intelligent environment needs to be an intelligent being [Dilger, 1998]. The environment must possess a degree of autonomy, adapt itself to changing conditions, and communicate with humans in a natural way. Second, adaptive software will be necessary for environment automation to become commonplace [Mozer, 1998]. It is unreasonable to assume the inhabitants will want to program the environment in order to customize the behavior, especially if the term "programming" implies traditional computer programming. Traditional approaches to user interfaces and software applications are not appropriate for intelligent environments.

This paper looks at intelligent environments from a networking and connectivity point of view in order to gain a fresh perspective. First, we define the qualities of an intelligent environment and present related research. Next, we provide a classification of the major components of an intelligent environment system. Then, we discuss Intelligent Internet Environments, and we describe several scenarios and associated challenges. Finally, we conclude with recommendations to guide future intelligent environment research.

2 CHARACTERISTICS OF INTELLIGENT ENVIRONMENTS

We begin by discussing common characteristics found in intelligent environments, such as *environment automation*, *natural interfaces*, *multimodal inputs*, *ubiquitous computing*, *aware computing* and *adaptable computing*. These qualities differentiate an intelligent environment from traditional environments, and are discussed below.

?? The goal of any intelligent environment is to automate the usage of the devices within the environment. The system automatically turns devices on and off, directs information to the appropriate devices, and simplifies the coordination of multiple devices. *Environment automation* allows an inhabitant to concentrate on the task at hand, rather than on the devices necessary to perform the task. It is important not to confuse centralized device control with environment automation. We do not consider an environment that only provides device control to be an intelligent environment. Automation is the key characteristic of an intelligent environment.

?? In general, computers do a poor job of supporting our daily activities, because they communicate with users via devices that are suitable to the computer. Unfortunately, the commonly used WIMP (windows, icons, menus, pointer) interface is exactly what limits the usefulness of a computer. The paradigm is not an acceptable form of human-computer interaction for future computing environments [Turk, 2000]. In an intelligent environment, Human-Computer Interfaces (HCI) should be *natural interfaces* designed for the human inhabitants and not for the computational system. Spoken language and human behavior should be used to communicate with the system, rather than pointing and clicking. A computational system should perform tasks in support of the human occupants naturally and non-intrusively. The goal of current HCI research is to remove the barriers of traditional computer interfaces. To date, almost all computer interface effort has concentrated on graphical user interfaces (GUIs). Work on the next generation of interfaces, namely Perceptual User Interfaces (PUIs), investigates techniques that seek to make natural interfaces by understanding how people interact [Turk, 2000].

?? The term *multimodal* is often used when discussing intelligent environments, because completely understanding an environment requires multiple sources of input. Multimodal interfaces can interact naturally with a user by using voice, hands, and even the entire body as communication devices. In a multimodal system, two or more sources of information are fused in order to reduce uncertainty [Oviatt, 2000]. An intelligent environment will be rich with devices, and thus must process multimodal information.

?? An intelligent environment requires *ubiquitous computing*, which is distributing the computation and network capability into the environment while also providing natural interfaces [Abowd, 1997]. Ubiquitous research concentrates on techniques that limit the intrusion on the physical environment. For example, furniture does not have to be "wired" in order to detect the presence of an occupant. Instead, vision and detection algorithms could be used to identify, locate, and track an individual. Ubiquity is the idea that the environment must be able to process information *as if* computational devices were everywhere, rather than actually embedding devices everywhere. It has been proposed that ubiquitous computing is the third wave of computing, following mainframe and desktop computing [Weiser, 1996].

?? A concept closely related to ubiquitous computing is *aware computing*. A system is provided knowledge about the user and environment in order to identify users, determine their focus of attention, and ascertain their intentions [Abowd, 1997]. The combination of ubiquitous and aware computing allows an intelligent environment to react to the desires of the inhabitants.

?? Finally, an intelligent environment must be *adaptive* since the world is a dynamic place. Environments evolve as the devices and people within the environment change. Even the physical layout of the environment may be altered over a period of time. The intelligent system must evolve with the environment.

Therefore, we conclude that ideal intelligent environments are adaptive systems that process multiple sources of input, interact naturally with the inhabitants in a ubiquitous manner, maintain awareness about the intentions of the inhabitants, and automate the environment. We emphasize, once again, that automation is the defining characteristic.

Current intelligent environment research projects do tackle the issues that arise when constructing environments with these qualities. In the following subsection, we summarize research efforts that emphasize understanding how to unobtrusively respond to the occupants' behavior in an omnipresent manner.

2.1 RELATED RESEARCH

One of the first intelligent environments was the Active Badge [Want, 1992] system installed in the Cambridge University Computer Laboratory. Personal identification badges were configured to emit unique infrared (IR) signals. Multiple sensor networks installed throughout the laboratory detected the emissions and relayed the information to standard workstations, which were part of an even larger network. The badge location was accurate to a 15-second time window. The system was so successful that the laboratory receptionist came to rely on the system to effectively locate personnel. Active Badge *automated* location activities by using *ubiquitous* resources.

Coen's Intelligent Room [Coen, 1998] is a HCI platform that uses computer vision and speech recognition systems. One configuration of the room supports disaster relief planning. Inhabitants can examine weather conditions by pointing to a map and asking for weather information of a geographic location. Since a major goal of the Intelligent Room is to make the user interface invisible, the room uses cameras and microphones to interpret the activities of the inhabitants of the room. Coen desires to enable unencumbered interaction with traditionally non-computational objects. The Intelligent Room incorporates *natural interfaces*, *multi-modal inputs*, *ubiquitous computing*, and *awareness* into its automated environment.

Mozer’s Neural Network House [Mozer, 1998] adapts temperature, lighting, and ventilation to the behavior and desires of the inhabitants. The neural network house required five miles of low-voltage conductor for collecting sensor data and a power-line communication system that controlled lighting, fans, and electrical outlets. The Adaptive Control of Home Environments (ACHE) system uses an instantaneous environment state to capture transformation and occupancy information. Decisions are split into two distinct activities. A setpoint generator determines a profile that specifies a target value, and a device regulator controls physical devices to achieve the desired goal. Two approaches to control are taken: indirect control using dynamic programming and models of the environment, or direct control using reinforcement learning. The system balances the goals of anticipating inhabitant needs and energy conservation. The Neural Network House automates the environment with *ubiquitous* and *adaptive* computing.

Brotherton and Abowd [Brotherton, 1998] have implemented the ultimate aid for university students, Classroom 2000. The classroom relieves students from the mundane tasks associated with lectures, such as note taking, so that they can focus on the lecture itself. Cameras, microphones, and projection technology are used to capture video, audio, and instructor notes. The information is provided to the students via the Internet. The classroom is designed to be unobtrusive, and a casual observer detects nothing too out of the ordinary when within the smart classroom. The environment records traditional classroom behaviors by passively monitoring instructor and student actions, and thus is *ubiquitous* and *aware*.

Finally, the Easy Living project [Hedberg, 2000] at Microsoft takes an approach similar to Coen's Intelligent Room. Fingerprints are used to identify inhabitants, and once identified their preferences are loaded. Cameras are used to locate users, and future plans include speech recognition. A room constructed for the project automatically continues a video at the point when an inhabitant last left the room. *Ubiquitous* and *aware* computing are characteristics of the room, with *natural interfaces* and *multimodal input* on the horizon.

Additional research projects and references can be found at our MavHome WWW site [MavHome, 2001]. These projects include the Intelligent Home Project at the University of Massachusetts, the Georgia Tech Aware Home, the Cisco Internet Home, the Verizon

Connected Family project, intelligent workspaces at Stanford and Xerox PARC, and the Nissan line of smart cars.

3 INTELLIGENT ENVIRONMENT SYSTEM

Before continuing, we will categorize and define the major components of an intelligent environment. In Figure 1, we identify two major categories. The *connectivity infrastructure* encompasses the physical hardware as well as required connectivity and communication software, while the *intelligence framework* identifies the software components necessary to provide automation of the environment.

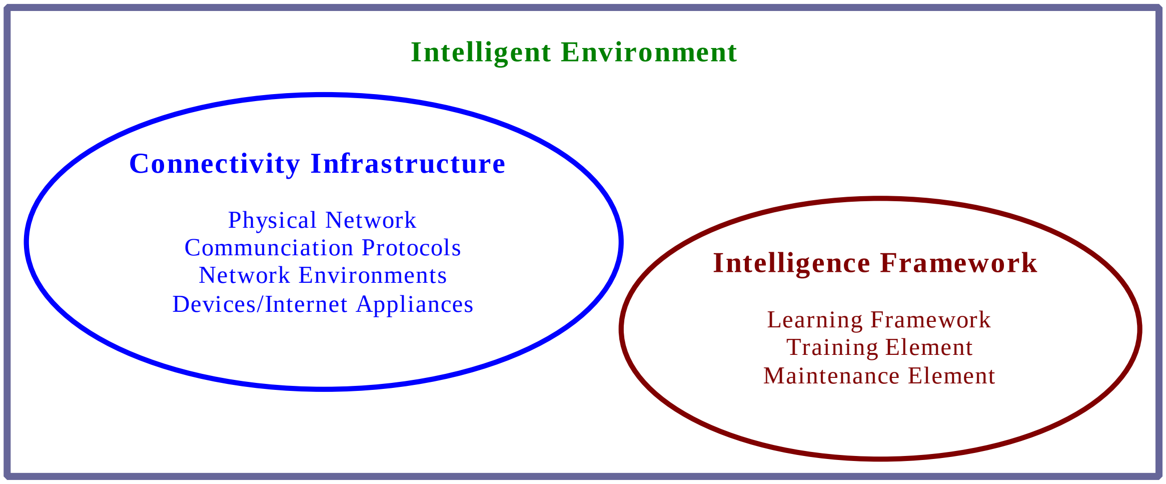

Figure 1 Intelligent Environment Infrastructure

This diagram is not intended to imply a particular implementation, relationship, or location of software and hardware components, but only to identify certain components and the capabilities each provides for the system. The only assumption made is that environment automation is provided. In addition, the diagram does not specifically address the other characteristics of natural interfaces, multimodal inputs, ubiquitous computing, aware computing and adaptable computing. The capabilities of the actual components incorporated into a system will determine what additional characteristics are available. Each of the components within the two major infrastructures is discussed next.

3.1 PHYSICAL NETWORK

Intelligent environment devices must be able to communicate with each other. Because the scope of activities performed by the devices varies, communication may take place over various mediums. A device may communicate on a local network contained within a house, or on a global network encompassing multiple local networks. Because of the proliferation of mobile devices, both wired and wireless network access will be necessary.

Multiple commercial approaches to the physical network either exist or are being explored. Power-line and phone line local networks can be constructed using the existing wiring in a house or building. Traditional Ethernet hardware can be used to provide high-end connectivity. New RF wireless technology, such as Bluetooth, can provide local wireless access. Global access can be achieved by using the Internet via gateways for wired connectivity, or by using the telecommunications infrastructure for wireless connectivity.

3.2 COMMUNICATION PROTOCOLS

The communication protocols for wired and wireless networks provide connectivity and communication rules for the devices. An important development is the merging of the telecommunication infrastructure and the Internet into what is often referred to as the next-generation network (NGN) [Moyer, 2001]. The Internet Protocol Version 6 (IPv6) specifically addresses mobility [Johnson, 2000], an issue normally addressed by telecommunication protocols. The next generation of cellular networks, 3rd Generation Cellular, will use packets and thus can support IP communication. The Session Initiation Protocol (SIP) [Schulzrinne, 2000], under development by the Internet Engineering Task Force (IETF), supports session based signaling, similar to Signaling System 7 (SS7) used by the Public Switched Telephone Network (PSTN). SIP provides personal mobility by using a single location-independent address identification to locate a user at more than one IP addresses. Finally, the Wireless Applications Protocol (WAP) adds Internet capabilities to devices with limited resources by the use of the Wireless Markup Language, which is derived from the eXtensible Markup Language (XML) [Stefan, 2000]. Internet capability can be extended to devices normally not part of the Internet. The ability to reach out and touch someone will soon include reaching out and touching a mobile device. Cellular and Internet compatible connectivity can be used by intelligent environment systems to extend the umbrella of coverage.

3.3 NETWORK ENVIRONMENT

An intelligent environment network must support dynamic collections of heterogeneous devices since devices will come and go as new ones are bought, old ones break down, and existing ones are moved around. In addition to signaling and communication, essential features of the network are device and service discovery, name resolution, and address binding. *Assumed knowledge*, *centralized knowledge*, and *decentralized knowledge* [Mittag, 2001] are three general approaches taken by network implementations. Examples of these approaches are provided below.

An *assumed knowledge* approach is characterized by prior knowledge of the device types and interfaces. PC developers dynamically connect applications on the same computer as well as across computers by specifying unique interfaces with approaches such as the Component Object Model (COM) and Distributed COM (DCOM). The Common Object Request Broker Architecture (CORBA) standard proposed by the Object Management Group (OMG) also supports communication between distributed objects/applications. The Aladdin project [Wang, 2000] takes advantage of COM in order to support extensible home networking applications. Dependability, extensibility, user-friendly interfaces and remote access capability are features of Aladdin. When an appliance attaches to an Aladdin network, a pre-defined device code specifies the device type. The network supports external devices, such as a cellular phone.

Jini is an example of the *centralized knowledge approach*, since a resource server maintains the service and device existence information. Jini, a distributed architecture sitting on top of Java, was developed to solve what can be described as the participant-to-participant network problem [Sun Microsystems, 2001]. Devices within one network, such as a cellular phone, request the services of a device in another network, such as the oven in a home. Traditionally, network components have interacted with known components distributed over the network. Participant-to-participant interaction requires the merging of networks. Services dynamically register traits to a Jini community, and the clients within the community discover and invoke the services of clients across the networks.

Finally, Universal Plug and Play (UPnP) represents a *decentralized knowledge approach* where resource information is distributed to the devices. Microsoft's Simple Control

Protocol (SCP) [Microsoft, 2001], designed to support communications mediums such as power-line communication, uses UPnP as the device model. A device advertises its services periodically, while client devices record the services and connect with the device when desiring to use a service. No single device acts as the controller.

3.4 INTELLIGENT DEVICES AND INTERNET APPLIANCES

The basic components of any intelligent environment are intelligent devices. One common perception of an intelligent device is any traditional device that contains an embedded processor that provides new functionality. If we use this definition as the minimum requirement for an intelligent device, then Internet Appliances can be considered intelligent devices.

The term Internet Appliance actually fits better than intelligent device, because the intelligence of many devices is fixed and contingent on the software installed at the time of construction. The term Internet Appliance can be used to cover a broad range of devices, from high-end devices with extensive computational power to low-end devices with very little computational power. Many such devices do exist.

When Internet Appliances for the home are discussed, the most common example is an intelligent refrigerator. In the future, our refrigerators may keep an inventory of their contents and automatically order replacements via the Internet. CHeF [Claymon, 2000] is a talking refrigerator at SRI International. It is connected to the Internet for online updates and monitors food inventory via electronic tags. Thus, CHeF can inform the user when the pantry is bare as well as the refrigerator. Vendors such as Frigidaire, Whirlpool, and Electrolux are also developing intelligent refrigerators. Other home devices under development include washing machines by Margherita 2000, vacuum cleaners by Electrolux, lawn mowers by Friendly Robotics, and even smart trashcans by NCR [MavHome, 2001].

3.5 LEARNING FRAMEWORK

The *learning framework* provides the intelligence backbone for the system, and three learning approaches can be used: *rote*, *supervised*, or *observation*.

?? With *rote* training, the user specifies the exact system behavior. Most commercial systems that provide centralized device control use the rote approach. For example, the user may program the VCR to record at a specific time, set the alarm clock for 7:30am, and turn the coffeepot on through the system. An environment using a rote approach could still be an intelligent environment. For example, if the system decides when the desired behavior is appropriate, then it is providing automation.

?? A *supervised* approach [Russell, 1995] requires the user to provide training examples, which are generalized by the system to predict future behavior. A mechanism is required that allows the user to indicate when supervised training is taking place.

?? Finally, with an *observation* approach [Russell, 1995] the system passively observes behavior and learns appropriate activities, thus requiring no direct interaction. A system using learning by observation is clearly the least intrusive. An observation approach is often referred to as *unsupervised learning*.

The amount of human intervention and associated implementation difficulty of each approach is depicted in Figure 2. The actual learning framework constructed for a system may incorporate more than one of the learning approaches. In addition, the learning approach may be distributed throughout the system by the use of intelligent agents.

At a minimum, the learning framework must provide automation. It may also support one of more of the other characteristics we have defined. Since the qualities of natural interfaces, ubiquitous computing, and aware computing imply less direct human intervention, we can conclude, by using Figure 2, that these qualities are also more difficult to implement.

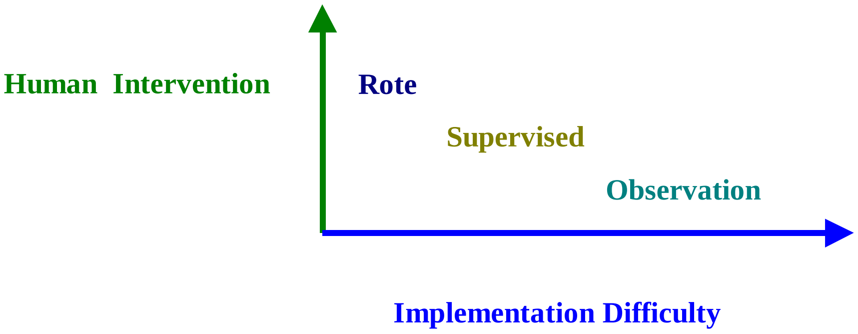

Figure 2 Learning Approaches

3.6 TRAINING ELEMENT

The learning approach determines the amount of effort needed for installation and training, which are tasks supported by the *training element*. With rote learning, the user must provide all desired behavior. With supervised learning, the user must indicate to the system when the training is taking place. With observation, the system decides when a sequence of actions is important. With all three approaches, it may be desirable to install the system with a generic default configuration.

3.7 MAINTENANCE ELEMENT

The *maintenance element* establishes how adaptive the system will be to environment changes. Rote learning requires the user to manually update desired system behavior. Supervised learning requires the user to provide new training examples and possibly even delete existing training examples in order to modify system behavior. It is the user's responsibility to recognize when maintenance is required with the two approaches, and there may be no difference between the training element and maintenance element. An observation approach requires no action on the user's part, rather the system must recognize when to adapt.

An important issue to consider is how the system will respond to drastic environment reconfigurations. For example, if a home is sold the new occupants will certainly have behavior patterns that differ from the previous occupants. The maintenance element determines how efficient and effective the system is at adapting to change.

4 INTERNET INTELLIGENT ENVIRONMENTS

Environments that consist of dynamically networked devices (Internet Appliances) capable of communicating with each other do exist. Research efforts continue on these components, and soon we will be able to connect a broader range of devices over a larger area of coverage in an increasing number of environments. The step of fitting common devices with computational and communicational capabilities is now in progress. Environments using these devices, however, tend to provide only centralized device control. The next step is to apply Artificial Intelligence techniques to provide *environment automation*. Rather than using traditional computer interfaces to manually configure and control Internet Appliances, we should intelligently automate the device activities and construct Internet Intelligent Environments.

4.1 SCENARIOS

The following scenarios describe what may be possible with Internet Appliance environments if an intelligence framework is added to the system.

4.1.1 Device Adaptation

An intelligent home environment is composed of common household devices such as a security system, a coffeepot, a television, an oven, and a heating system. The device configuration changes as the homeowner buys, upgrades, and discards household appliances. The goal of the system is to automate the tasks manually performed by the occupants with these devices, so the system continually adapts to the dynamic device configuration. The connectivity infrastructure reconfigures and seamlessly manages the devices found, while the intelligence framework continually adapts the system behavior to incorporate the configuration changes.

4.1.2 Usage of Internet and Telecommunication Networks

The Internet contains a wealth of information. An Internet-capable intelligent environment uses the WWW to upgrade system and device capabilities by downloading upgrades. The environment monitors device capabilities and compares current capabilities with those found in WWW-based upgrades. A download is performed when the system determines that improvements are desirable or necessary. The very term “Internet Appliances” implies Internet connectivity, and thus the ability to upgrade autonomously is possible.

A telecommunication-capable intelligent environment uses the wireless infrastructure to coordinate the activities of mobile devices. Communicating with devices that are not physically part of the environment extends the umbrella of coverage of the system. An intelligent home system that uses the wireless digital network can extend anywhere wireless devices can be carried. If you leave the iron on when traveling on a family vacation, the home shuts it off and forwards a message to your cell phone so you don’t worry if you turned it off or not. Providing automation to these networks increases their capabilities.

4.1.3 Temporal and Contextual Behavior

The lifecycle of a room is very dynamic. Consider a game room located in a household with adults and children, and the room contains a television and VCR. The kids like to watch cartoons on Saturday mornings. As the kids get older, their taste in cartoons shifts from instructional toddler shows to action–based cartoons.

Dad, on the other hand, likes to watch baseball games in the evenings and on weekends. Unfortunately, his favorite team is not doing too well, so after July he stops watching the games. Dad is waiting till next April when the season starts again. Of course, baseball is on hold in April and May if the hometown hockey team is winning.

This year, Mom is really excited about cooking. Last fall it was stock market investing and last summer it was painting with watercolors. It seems Mom’s interests depend somewhat on the current activities of her professional women’s club.

This game room system accounts for individual behaviors of the inhabitants by making continual adjustments that optimize the program selections of the inhabitants. The appropriate cartoons are on when the kids wake up on Saturday. The game is taped for Dad when he is unexpectedly out of town on business. When Mom has a free moment, a cooking show is offered for her viewing pleasure. Adaptation and awareness are the defining characteristics of the system.

4.2 MAJOR CHALLENGES

The above scenarios illuminate the challenges that arise with intelligent environments, and the factors that contribute to the lack of actual intelligent environments. Some of the obstacles that would be encountered by the scenarios, as well as current research projects, are discussed here.

4.2.1 Inherent Difficulties

One reason we do not find an abundance of intelligent environments is that attempting to automate activities remains difficult. The characteristics we defined for intelligent environments - environment automation, natural interfaces, multimodal inputs, ubiquitous computing, aware computing and adaptable computing - cannot be provided by simple solutions. As we pointed out previously, the intelligence framework required to support these qualities is difficult to implement. Significant effort and ingenuity are required to field implementations for essential components.

As an example, consider the natural interfaces desired for intelligent environments. Franklin [Franklin, 1998] bluntly states that natural interaction is all but absent from computer systems, because it is really hard to do. Computer interfaces are easier to develop, so human-computer interaction is tailored for the computer and not the user.

4.2.2 Specialization

Inhabited environments are unique. Rarely will two environments consist of the same physical objects, physical dimensions, or human occupants. Many current research projects, however, use a fixed environment with predefined devices. An important question to ask is: can the system adapt itself to activities and devices not previously considered? And, what kind of human intervention is required for adaptation? It is impossible to widely deploy an intelligent system that is specific to a narrowly defined environment.

The paradox, however, is that research projects are often very specific in nature because of the difficult problems being addressed. The environment contains predefined devices and software in order to investigate particular issues. Few research projects address the notion of a system that does not have complete a priori knowledge of the devices and their intended use.

4.2.3 Heterogeneity and Intrusiveness

A typical environment is composed of devices from numerous sources. Look at the appliances in a home and you find devices that have been built by various companies with varying capabilities. The configuration is dynamic, because devices are added, upgraded, and thrown away. In the near future, the capability exists for a device to be upgraded via the Internet. The device you have today may not be the same tomorrow.

Cameras, microphones, miles of wiring, and even robots are some of the components required by current intelligent environments. Although such components could be incorporated into newly constructed environments, they are intrusive and expensive to install into existing environments. In addition, not all of these components may be readily acceptable by potential inhabitants, especially in a home environment. The least intrusive intelligent environment systems are those that can support common devices, such as Internet Appliances.

4.2.4 Temporal (Contextual) Sensitivity

Handling time, synchronicity, and streams of events in a dynamic environment is a difficult sub-problem of an intelligent environment. Eric Horvitz [Horvitz, 1996] states “we must endow systems with the ability to represent and reason about the time-dependent dynamics of belief and action, including such critical notations as the persistence and dynamics of world states.” Intelligent environments must process data and behaviors that are temporal in nature and context sensitive.

It is desirable for a computer to understand the context of a person’s action, consider the goals of a person, and understand what responses are appropriate to different goals and sequences of actions [Franklin, 1998]. Unfortunately, this level of understanding and action taking is a very difficult problem to solve, but is necessary for the system to be adaptive and aware.

5 RECOMMENDATIONS

Although significant challenges remain in the area of intelligent environments, the connectivity infrastructure components have advanced quite rapidly over the past several years. We assert that the recent advances made in the connectivity infrastructure will aid research efforts on the intelligence framework components. Easily constructed environments will help further automation research by providing needed test platforms. It has been pointed out that a common platform that can support thousands of distributed objects could greatly aid the intelligent environment research community [Hasha, 1999]. We also assert that existing environments can be improved by conducting research focused on adding an intelligence framework to common connectivity infrastructures. Soon, an abundance of Internet Appliances and Internet-capable platforms will be readily available, but environment automation will be missing in most of these environments. We therefore make the following recommendations, which aim to take advantage of the next generation of networked environments in order to produce Internet Intelligent Environments.

5.1 ADAPTING TO DYNAMIC COLLECTIONS OF DEVICES

Emerging network technology can support dynamic collections of devices. As we discussed previously, device and service discovery are properties of several network approaches. Thus, the connectivity infrastructure of an environment will be able to support dynamic device configurations and can serve as a test platform for research on intelligently controlling dynamic device collections. Changing the configuration of the network may be as easy as plugging and unplugging devices from an outlet. We mentioned that research is lacking on environments with dynamic configurations, such as the home. The next generation of networking will provide the platforms necessary to explore intelligence frameworks that can automate device usage in a dynamic environment, which is one of the difficult challenges with intelligent environments.

5.2 AUTOMATING THE USAGE OF SIMPLE DEVICES

Many intelligent environments use multimodal and natural interfaces that require the integration of a wide array of Artificial Intelligence subdisciplines. These projects focus on advanced devices such as speech, vision, and hearing, and the difficult problems they must address can be viewed as AI-Complete problems [MavHome, 2001] when evaluated in terms of problem complexity.

On the other hand, traditional home appliances, which are now evolving into Internet Appliances, can also provide multimodal input. The devices may not provide as much information as more advanced devices, yet they are much more likely to be found in an environment. Some desired qualities, such as natural interfaces, may be missing, but the key goal of automation can still be achieved.

By taking advantage of commonly used and familiar devices to obtain environment inputs, the system may be less expensive and gain greater acceptance. Home automation could benefit significantly from research looking at collections of relatively simple devices.

5.3 LEARNING OCCUPANT BEHAVIOR

Since a common goal of most intelligent environments is reducing human actions, it would seem that analyzing the occupants is one way to automate human actions. In many environments, the devices may be fixed but the user intent is different A conference room in a corporation may contain the same devices as a classroom at a university. A single classroom in a university may be used by a variety of instructors teaching a numerous courses to a wide spectrum of students. It is the occupant behavior that will differ in these environments. An intelligence framework that analyzes occupant behavior is useful to multiple environments, rather than just one specific environment. Research activities that focus on learning and predicting user actions, such as predicting UNIX commands [Davison, 1998], may be applicable. Learning and predicting occupant behavior would reduce the specialization of an intelligent environment system.

5.4 ALLOWING USER INTERVENTION

We have discussed how the implementation difficulty increases as human intervention decreases. However, a small amount of human intervention can help overcome difficult problems in an environment [Brotherton, 1998]. Therefore, it may be acceptable to have the occupants use non-traditional programming to help shape the behavior of the environment, even though it is not desirable to have the occupants of an environment perform traditional programming tasks. Given that only a fraction of a computer's power will be accessible unless nonprogrammers can program the computer [Smith, 2000], some intervention seems unavoidable.

One possible approach is Programming By Example (PBE) [Lieberman, 2000], which can be used to overcome a brittle pre-determined infrastructure. Theoretically, the approach allows a novice to program a complex system. With PBE, the user demonstrates desired behavior to a software agent. The behavior is generalized by the agent and recorded as a program to be executed later. The user instructs the software agent.

It may be possible to use PBE to allow a user to program their intelligent environment. The outcome of the learning may be the creation of agents that control the interaction of various individual agents/devices. Up to now, these are agents have been created statically prior to the installation of the system. In addition, learning is not limited only to active learning. Both instructible interface agents (behavior based) and passive (knowledge based) agents that gather knowledge incrementally [Lieberman, 1996] are possible.

Consider Lieberman's ALIVE project [Lieberman, 1996], which uses a virtual playmate to aid a user in the completion of simple tasks. The playmate observes the user and pitches in when it determines something it can perform for the user. The same approach could be used as part of the training element and even maintenance element components of an intelligence framework. An approach like PBE could help conquer several of the inherent difficulties we have outlined with intelligent environments, by allowing acceptable human intervention.

5.5 LEARNING TEMPORAL PATTERNS

One of the most complex problems to solve is the temporal and context-sensitive nature of intelligent environments. The usage of the television provided in the game room scenario is very temporal and context dependent. The temporal frames may be days, weeks, months, seasons, years, even decades. The context varies depending on the person and may also be temporal. Learning to process and respond to temporal information is important for system awareness. Internet Appliances are devices in which the usage is very temporal and contextual in nature. Learning and planning techniques can be used to understand home behavior, inhabitant behavior, and infer the results of selected actions [MavHome, 2001], and an environment constructed with Internet Appliances can provide an inexpensive test platform for this research.

6 CONCLUSIONS

Because of the progress in Internet Appliances, communication protocols, network environments, and network components, the intelligent environment field is rich with challenges and ripe for research. Existing device environments could benefit from research devoted to device automation, and the connectivity infrastructure components could improve intelligence framework research by providing easily constructed device environments. We intend to investigate intelligent environment challenges with our MavHome project [MavHome, 2001], using the above recommendations to guide our research. By addressing the challenges we have outlined, the home automation environments that were predicted long ago may finally emerge in the form of Internet Intelligent Environments.

7 REFERENCES

1. G. Abowd and I. Essa, “Ubiquitous and aware computing,” http://www.cc.gatech.edu/fce/pubs/horizons.html, 1997.
2. J. Brotherton and G. Abowd, “Rooms Take Note: Room Takes Notes!,” *AAAI Spring Symposium*, Stanford, pp. 23-30, March 1998.
3. D. Claymon, “Future PCs will fade into background, scientists say,” *Dallas Morning News*, section F, p. 7, January 13, 2000.
4. M. Coen, “Design principles for Intelligent Environments”, *AAAI Spring Symposium*, Stanford, pp. 36-43, March 1998.
5. B. Davison and H. Hirsh, "Predicting Sequences of User Actions," *Technical Report*, Rutgers, The State University of New York, 1998.
6. W. Dilger, “A society of self-organizing agents in the Intelligent Home,” *AAAI Spring Symposium*, Stanford, pp. 55-59, March 1998.
7. D. Franklin and J. Flachsbart, “All gadget and no representation makes Jack a dull environment,” *AAAI Spring Symposium*, Stanford, pp. 155-160, March 1998.
8. R. Hasha, "Needed: A common distributed-object platform", *IEEE Intelligent Systems*, vol. 14, no. 2, March/April 1999.
9. S. Hedberg, "After desktop computing: A progress report on smart environments research", *IEEE Intelligent Systems*, vol. 15, no. 5, pp. 7-9. September/October 2000.
10. E. Horvitz, et al., "Challenging problems for Artificial Intelligence," AAAI-1996.
11. D. Johnson and C. Perkins, *Internet Draft - Mobility Support in IPv6. http://www.ietf.org/internet-drafts/draft-ietf-mobileip-ipv6-13txt*. November 2000, work in progress.
12. H. Lieberman, “Programming By Example,” *Communications of the ACM*, vol. 43, no. 3, pp. 72-74, March 2000.
13. H. Lieberman and D. Maulsby, “Instructible agents: Software that just keeps getting better,” IBM Systems Journal, vol. 35, no. 3&4, 1996.
14. "MavHome: An Intelligent Home Environment," http://www-cse.uta.edu/smarthome, University of Texas-Arlington, 2001.
15. Microsoft Corporation, "Simple Control Protocol Vision", http://www.microsoft.com/homenet/scp, Microsoft Research, Redmond, WA, 2001.
16. L. Mittag, "Internet device Interaction Control," *Embedded Systems Programming*, vol. 14, no. 2, pp. 63-70, February 2001.

17. S. Moyer and A. Umar, "The impact of network convergence on telecommunications software", *IEEE Communications Magazine*, vol. 39, no. 1, pp. 78-84, January 2001.
18. M. Mozer, “The Neural Network House: An Environment that Adapts to its Inhabitants”, *AAAI Spring Symposium*, Stanford, pp. 110-114, March 1998.
19. S. Oviatt and P. Cohen, “Multimodal Interfaces that process what comes naturally”, *Communications of the ACM*, vol. 43, no. 3, pp. 45 – 53, March 2000.
20. S. Russell and P. Norvig, "Artificial Intelligence - A Modern Approach," Prentice-Hall, 1995.
21. H. Schulzrinne and J. Rosenberg, "The Session Initiation Protocol: Internet-Centric Signaling," *IEEE Communications Magazine*, vol. 38, no. 10, pp. 134-141, October 2000.
22. D. Smith, A. Cypher, and A. Telser, “Novice programming comes of age,” *Communications of the ACM*, vol. 43, no. 3, pp. 75 – 81, March 2000.
23. J. Stefan, "Working with WAP", *Embedded Systems Magazine*, vol. 13, no. 11, pp. 47-48, October 2000.
24. Sun Microsystems, "Jini Network Technology: An executive overview", *Technical Report*, Sun Microsystems, Inc., Palo Alto, Ca, 2001.
25. M. Turk and G. Robertson, "Perceptual User Interfaces," *Communications of the ACM,* vol. 43, no. 3, pp. 33-34, March 2000.
26. Y. Wang, W. Russell, and A. Arora, "A toolkit for building dependable and extensible home networking algorithms," *Technical Report MSR-TR-2000-07*, Microsoft Research, Redmond, WA, February 2000.
27. R. Want, A. Hopper, V. Falcao, and J. Gibbons, "The Active Badge location system," *ACM Transaction on Information Systems*, vol. 10, no. 1, pp. 91-102, January 1992.
28. M. Weiser, "Ubiquitous Computing", http://www.ubiquitous.com/hypertext/weiser/UbiHome.html, 1996.