At Home with Ubiquitous Computing: Seven Challenges

W. Keith Edwards and Rebecca E. Grinter

Computer Science Laboratory, Xerox Palo Alto Research Center
Palo Alto, California 94304, USA
{kedwards,grinter}@parc.xerox.com

Abstract. The smart home offers a new opportunity to augment people's lives with ubiquitous computing technology that provides increased communications, awareness, and functionality. Recently, a number of trends have increased the likelihood that the aware home can soon become a reality. We examine a number of challenges from the technical, social, and pragmatic domains that we feel must be overcome before the vision of the smart home, posited by ubiquitous computing research, can become a reality. Our hope in raising these issues is to create a conversation among researchers in the varied disciplines that make up ubiquitous computing. In particular, we hope to raise awareness of the existing literature on the adoption, use, and history of domestic technologies, as well as the use of situated studies, and the benefits that these can bring to bear on the design and evaluation of technologies for the home

Keywords. Home, ubiquitous computing, context-awareness, domestic technologies, evaluation

1 Introduction

The notion that we could eventually live in so-called "smart homes"—domestic environments in which we are surrounded by interconnected technologies that are, more or less, responsive to our presence and actions—seems increasing plausible. Trends such as Moore's Law, the proliferation of networkable devices, wireless technologies, and an increasing vendor focus on technologies for the home (perhaps arising from a belief that the office is a technology-saturated market) are driving awareness of the smart home idea out of academia and into mainstream thinking [5, 12]. Still, however, the most advanced glimpses of the potential future of domestic technologies can be found in settings such as the Aware Home laboratory at Georgia Tech [10].

While initiatives such as the Aware Home can teach us about what a smart home could provide to its "occupant-users," social and technical questions remain. In this paper we present a number of challenges that we feel must be overcome before the smart home concept can begin to move to reality. These challenges are based on field studies, research in progress, and analysis of the adoption of other domestic technologies. These challenges concern not just technical direction, but also the social and

G. D. Abowd, B. Brumitt, S. A. N. Shafer (Eds.): Ubicomp 2001, LNCS 2201, pp. 256-272, 2001.

ethical directions of ubiquitous computing in the home. Our purpose in this paper is two-fold. First, we hope to illuminate some implications of technical change that must be understood in order to produce domestic computing technology that is not simply ubiquitous, but also calm [28, 29]. Second, we hope to show that field studies and the existing literature on the adoption, use, and history of domestic technologies can support and influence the design and evaluation of technologies for the smart home.

2 The Seven Challenges

In this section we present the seven challenges; these challenges are based in the technical, social, and pragmatic domains. They cover problems that arise from the ways in which we expect smart homes to be deployed and inhabited; technical questions of interoperability, manageability, and reliability; social concerns about the adoption of domestic technologies and the implications of such technologies; and design issues that arise from considering just how smart the smart home must be.

2.1 Challenge One: The "Accidentally" Smart Home

Current research into domestic technologies has, for obvious reasons, taken an intentional approach to designing and building the smart home. That is, the environment has been designed from the ground up to support and evaluate technologies deployed there. Since pervasive infrastructure for ubiquitous computing does not exist in today's homes, houses must be explicitly outfitted for these sorts of technologies, and new applications must be created that are specifically written to serve as a test bed for smart home research.

However, while new homes may eventually be purpose-built for such smart applications, existing homes were not designed as such. Perhaps homeowners may decide to "upgrade" their homes to support these new technologies. But it seems more likely that new technologies will be brought piecemeal into the home; unlike the "lab houses" that serve as experiments in domestic technology today, these homes will not be custom designed from the start to accommodate and integrate these technologies. We call this phenomenon the "accidentally" smart home—a home that contains an accretion of technological components embedded in an environment that has not benefited from a holistic, ground-up approach to design and integration.

Ignoring for a moment the implications when disparate (and potentially conflicting) technology is meant to interpret and act on the behavior of its occupants[1], we can envision plausible problems with even the most mundane smart home technologies. Imagine, for example, that homeowners wake one weekend, and come downstairs looking forward to their first cup of coffee and Cartalk on NPR.[2] To their surprise, no sound emerges from their speakers. The reason is that their neighbors have purchased new Bluetooth-enabled speakers which, when first installed, associate themselves

[1] See challenge seven.
[2] Cartalk did not sponsor this research, it's just a very good show.

with the nearest sound source; in this case, the original homeowners' Bluetooth-enabled stereo. Meanwhile, the neighbors are abruptly awakened to tales of car repair.

This is one scenario made likely by the accidentally smart home. The general question that this scenario raises is how will the occupant-users adapt to the idea that their home has suddenly reached a level of complexity at which it becomes unpredictable. A specific question this scenario raises is how will they begin the process of making sense of what has happened. How will they begin the process of "debugging" their home to determine what has happened to their speakers?

In this simple scenario, it is precisely the wirelessness of the speakers that makes them compelling to the homeowners—the speakers can be untethered and placed exactly where they want to hear music. But it is also this same wirelessness that does away with the traditional affordances for understanding the connectivity between the speakers and the stereo (or indeed that the speakers can and should be connected to the stereo in the first place).

In particular, the homeowners may not realize that their wireless speakers can actually connect themselves to sound sources in another house as easily as to sound sources within the home. Models of connectivity are explicit when physical wires are present: the "range" of connectivity is apparent, connections are observable, and connections don't change on their own. The intangible models of connectivity that wireless technologies bring must be learned.

The general question, then, is how will occupant-users build up a model of how to control, use, and debug technologies that will interact with one another in the environment? What will the experience of the home as a whole be when these technologies are brought in gradually, and without the benefit of a top-to-bottom design? Will the occupant-users be prepared to manage their smart home when the time comes? Particularly when these complex technologies offer fewer physical affordances than we are used to?[3]

Technology positivists may say that this is "simply a design problem," and they are correct in the sense that the underlying technology itself does not dictate the behavior described here. Perhaps future models of connection will require that homeowners set a security key for all of their devices, or the vendor of the neighbors' speakers will develop a UI so intuitive and reliable that the neighbors would never make this mistake.

While this may be true, such a sidestep doesn't remove the fact that the situation described above *does* present a complex design challenge. The design challenge is to provide affordances to help users understand the technology. Consider, for example, the recent publicity given to IEEE802.11b wireless networking security failures. This technology—with industry backing and emphasis on ease of use and security—has famously provided impromptu connections to law firms, corporate offices, and development houses from passers-by in cars an on park benches [3, 18]. Here is an exam-

[3] While the focus in this section has been on intelligibility problems that arise when there is an accretion of interoperable technology, we believe that these problems are inherent in many of the visions of ubiquitous computing, in which technological artifacts, computational processing, and environmental sensing are made—to some degree—invisible and inaccessible to their users.

ple of a real world "design problem" much like the one described above, and which has clearly not yet been solved. Concisely, the problem is one of intelligibility in the face of radical—and perhaps unexpected—connectivity.

If we take as a given that few homes in the "real world" will ever be designed, top to bottom, as a holistic system of well-meshed, interoperable components, then a number of questions become important.

- What kinds of affordances do we need to provide to occupant-users to make the system intelligible? (e.g., Is the device recording, displaying, manipulating information about me)
- How can I tell how my devices are interacting? (e.g., What are my devices interacting with, and how do they choose?)
- What are the boundaries of my smart home? (e.g., What are the walls? How much privacy do I have?)
- What are the potential configurations of my devices? (e.g., What connects with what, what won't connect, and why?)
- How can users be made aware of the affordances of the entire home itself? (e.g., what are the possible and impossible configurations of this home?)
- Where will the locus of interaction be in a system that exists in no one place, but rather represents the sum of many interoperable (and changing) parts? (e.g., where does the UI *live*?)
- How do I control these devices, and the whole system? (e.g., Where are the controls, what visualizations of the whole system do I have?)

Current domestic technologies—with their only limited ability to connect with one another, and strong affordances of connection—do not provide good models for the smart home. Such a home will need to present its occupants with an intuitive sense of the possibilities it affords, the current state of the systems within the home, interfaces for controlling the systems in the home as a whole, and a means by which "accidents" (such as a neighbor hijacking their speakers) can be repaired or—even better—prevented in the first place. And these abilities must be provided and maintained in an environment in which new devices are added, old devices are removed, devices from different manufacturers may coexist, and wireless connectivity may extend beyond the walls of the home itself.

The challenge for homeowners with these devices will be to understand when their houses make the transition from dumb to smart and manage that transformation. The challenge for ubiquitous computing is to help homeowners understand their accidentally smart homes by providing insights into what these devices can do, what they have done, and how we control it.

2.2 Challenge Two: Impromptu Interoperability

The previous section discussed the challenge of ensuring that an environment will be intelligible when it comprises a number of components, each of which may have been acquired at different times, from different vendors, and which were created under different design constraints and considerations. And yet while this is clearly a crucially important challenge, it is predicated on the notion that such disparate components will be able to interoperate *at all*.

We believe that impromptu interoperability—not just the simple ability to interconnect, but the ability to do so with little or no advance planning or implementation—is implicit in much of the current literature of ubiquitous computing in the home. With fluid, impromptu interoperability, individual technologies have the potential to create a fabric of complementary functionality. Without it, the smart home of the future is likely to be characterized by islands of functionality, as the sets of devices that were explicitly built to recognize each other can interoperate, but other sets of devices cannot. (Such a world is likely to be one of software upgrades, version mismatches, and driver installations, which leads to our third challenge in the next section.) Such interoperability, while a challenge in its own right, increases the challenges of intelligibility as discussed in the previous section.

The chief obstacle limiting such impromptu interoperability now is that, in general, every device or software service must be explicitly written to understand *every other type* of device or software that it may encounter. If the applications on my PDA are to be able to print, then those applications (and the operating system on which it is built) must be explicitly written to understand and use the notion of a "printer"—what such a thing is, how to communicate with it, and why one would talk to it in the first place.

Without this *a priori* agreement on both syntax and semantics, interoperability is difficult if not impossible. And yet the smart home (as well as other visions of ubiquitous computing outside the home) posits the existence of a rich fabric of devices and software, somehow all seamlessly interconnecting with one another.

Must we agree on a complete set of standards for how these entities will be defined and used, known to all parties before any implementation can begin? Will we have to restrict our environments to only using devices and software that "fit" with the protocols already in place?

This challenge goes beyond mere standards. While standards for particular domains—printing, image capture, data storage—allow an entity to communicate with an entire *class* of devices or services using a standard protocol, they do not alleviate the core problem: that it is implausible to expect that all classes of devices or services will be known to all others, and that we can thus define standards for every type of device or service *a priori*. Instead, new models of connectivity are needed.

Research has begun to explore such models. Most of these models work by standardizing communication at the syntactical level (protocols and interfaces) and leaving to it a human to impose semantics. The event heap work at Stanford [9], for example, establishes a common tuple space protocol that all parties agree to implement. Particular tuples in the space may have meaning to certain parties, however, and that semantic agreement is implemented by the developer ("this tuple represents a request to scan an image," for example). The CoolTown project [16] leverages existing protocols (HTTP) and content encodings (HTML) to allow arbitrary entities that "understand" the language of the web to interact. The Speakeasy project at Xerox PARC [7] defines a set of interfaces that leverage mobile code to extend the behavior of entities in the environment; the end user provides the semantic knowledge to decide when and whether to use a particular entity.

These projects represent steps toward new models of interconnectivity, but the problem space is large and novel. Our challenge is to ensure that the future of the smart home is not one of incompatibility and isolated islands of functionality, but rather one in which occupant-users can expect the systems in their home to work together fluidly. We believe that this challenge requires radical new models of connectivity and interoperability that reach beyond simple prior agreement on standard protocols and interfaces.

2.3 Challenge Three: No Systems Administrator

As computers enter the home in greater numbers, individuals find themselves becoming systems administrators. Indeed, the average home computer user now has to concern herself with chores that would seem familiar to a mainframe systems operator from the days of the high priesthood: upgrading hardware, performing software installation and removal, and so on. The advent of always-on broadband connections and in-house networks have finally brought to our homes the few systems administration tasks that had so far eluded us: network and security administration. These are chores that are overwhelmingly complex and understood by few, even among "early adopters."[4] What will the situation be when our homes are filled by complex technological artifacts that are meant to interoperate with each other and with the outside world?

As designers of technology, we cannot plausibly expect such advanced knowledge of potential occupant-users of the smart home, if we expect anyone to actually wish to inhabit such homes. Indeed, if the lack of ability or interest in home "administration" chores as mundane as plumbing, electrical wiring, or appliance repair is any indication, there will effectively be *no* systems administrator in the smart home.

How, then, will we design technologies for the smart home that require no on-site expert? Fortunately, there are models for administration-free use of complex technologies other than general-purpose computing systems.

Traditional appliances, for example, are single-function devices that provide simple controls, straightforward affordances, and generally good ease of use (most people can use the office microwave oven without reading the instruction manual, for instance). When such a device breaks (which happens rarely), users are not expected to fix it themselves. Instead, an expert is called who comes to the house to make the repair.

There has been a move, recently, toward "appliance-centric" computing in which digital devices embody some single function [19]; we consider it an open question, however, as to how well this approach will scale, especially when such appliances are asked to interact with other sorts of devices in fluid ways (see Challenge Two).

Perhaps a more fitting model for administration in the smart home can be found in existing utilities, such as the telephone and cable television networks. In the utility model, most of the "intelligence" in the system resides in the network itself. The

[4] Note, for example, the recent publicity given to home computers being systematically hacked via their persistent broadband connections, a problem once confined to large companies, governments, and universities [13]. Consequently, sales of home firewalls are surging [23].

home contains only the most simple and minimal "front end" functionality needed to access the network. The telephone system is, of course, the most well-known example of this model: a simple, rotary telephone can be used to access any other telephone in the world, including cellular telephones that didn't exist at the time the rotary phone was built. This expanding functionality is available because the sophistication of the back-end network is increasing. The cable TV network, with its set-top boxes, is another example of the utility model, as are ISPs such as AOL and MSN, who bundle and preconfigure their networking software to create turnkey internet access points.

Generalizations of this model have been proposed by others as a solution for "outsourced" home administration, by organizations such as the Open Services Gateway Initiative [21].

Either of these approaches—the appliance model, or the utility model—brings with it a number of attendant technical and design challenges. In the appliance model, the challenges are largely in the design domain: how can these small devices deliver rich interactions with an ever-expanding coterie of technology in the home, without losing the simplicity that its their raison d'etre? In the utility model, how can we design technical solutions for remote diagnosis, administration, and software upgrades (in particular, with the security to prevent the kid next door from performing his own, unwarranted, remote diagnosis, administration, and upgrades)?

Regardless of the overall model chosen, occupant-users will still have some administration that they will have to do, simply because not all of the dynamics of the home can be known by the developer of the appliance, or the owner of the utility. The particular ways in which individual devices are used by members of the home, for example, may need to be reflected in configurations, security parameters, and device interactions that can only be implemented by the owners of those devices—not some external third party.

These issues of domestic technology usage form our fourth challenge.

2.4 Challenge Four: Designing for Domestic Use

The last decade has seen many studies showing how users adopt technology in surprising and unpredictable ways. Most of these have focused on office technologies, which are notably different than technologies for the home [15]. We agree with Abowd and Mynatt [2] that there is a need for studies of domestic settings to inform design. In this section we argue that studies of how the telephone and electricity were adopted in the home provide compelling evidence for further studies, and show what studies of the modern home reveal about design of technologies.

The telephone must be among the most ubiquitous technologies in the home. The study of its adoption reveals that while its inventors foresaw a social role for the phone, its initial vendors did not [8]. The telephone company did not believe that sociability was an important or appropriate use of their technology. It was not for several decades, and after the telephone was broadly adopted, that the Bell System promoted the device as a mechanism for having conversations with distant friends and family.

The adoption of the landline telephone could be viewed as a triumph of user persistence over vendor beliefs. Recently, phone adoption has received new attention because of wireless devices. Palen et al. [22] observed that individuals tend to purchase wireless phones for emergency and coordination reasons, and do not consider sociability to be important. However, within weeks of purchasing the phone these same owners used it for social calls.

The adoption of landline and wireless phones suggests that vendors and even users find it hard to foresee how they will use a technology. Electricity, another pervasive domestic technology, shows that new uses sometimes do not last. At the turn of the century, the homes of the wealthy were often outfitted with electrically-conducting rails in the floors; "electricity girls," equipped with metal shoes and wearable light fixtures, would entertain party guests by moving from room to room, carrying their own illumination [20]. Findings from these analyses reinforce the need for conducting studies of domestic settings and relying on analysis of the stable and compelling routines of the home, rather than supposition, company dictate, fad, or marketing.

Recent studies of domestic settings have taken this approach. They highlight a variety of findings, many of which stem from the fact that domestic technologies are not "owned" by an individual. Many are governed by household rules that determine: who uses what device, when, where, whether they pay, how old they are, and for what purposes.

For example, in their study of set-top box use in various homes, Hughes et al. [14] describe a relationship between technology use and space "ownership" within the home. They observed that occupants used technologies such as the television to indicate that they controlled behavior in that part of the home. They found that others knew and respected these routines. When occupants had conflicts over television use, they settled disputes by buying another television or making the current one more mobile. Finally, they observed that the television accommodates multiple usage requirements by making it possible for different occupants to watch their own programs [14]. Video and TiVo technologies make the television even more accommodating. Television and its associated technologies fit into the home by being portable and flexible to occupants' requirements.

Our study of wireless text messaging in the home shows how devices are used and shared. We found that the teenagers used text messages to arrange times to talk on the landline phone or use the computer to Instant Message [11]. Since both the phone and the computer were shared devices in their own homes and their friends' houses, teens used a technology that they individually owned to coordinate times when they all had access to those shared devices.

We also found that teenagers used "quiet" technologies such as text messaging to avoid disturbing the routines of other people. Quiet technologies do not ring or require voice interactions. Text messaging was quiet, and consequently allowed the teenagers to communicate without other household members being aware of or disturbed by the interaction. In this case text messaging meets the requirements of its users as well as those who are not using it but are sharing the same space.

In summary, smart technologies—indeed *any* technologies—will be disruptive to the home environment. Predicting these disruptions is difficult, as illustrated by the cases at the opening of this section. The challenge for designers, then, is to pay heed to the stable and compelling routines of the home, rather than external factors, including the abilities of the technology itself. These routines are subtle, complex, and ill-articulated, if they are articulated at all; thus, there is a great need for further studies of how home occupants appropriate and adapt new technologies. Only by grounding our designs in such realities of the home will we have a better chance to minimize, or at least predict, the effects of our technologies.

2.5 Challenge Five: Social Implications of Aware Home Technologies

Understanding how technologies fit into daily routines is one aspect of designing the smart home. However, technologies have other social implications that also bear examination. In this section, we describe some social implications of aware home technologies that merit discussion within this community, as well as presenting opportunities to engage other research disciplines in discourse about the future we are designing.

Abowd and Mynatt [2] have addressed the social implications of ubiquitous computing, privacy in particular. We believe this focus is very appropriate, since privacy is important. However, we believe that there are other broad social implications of domestic technologies which are not as widely explored by members of the ubiquitous computing community. Studies illustrate other potential consequences of domestic technologies, and we focus on two of these: "labor saving" and good parenting.

Some historical studies have challenged the belief of technologies as being labor saving devices. The washing machine is one of those technologies. The washing machine was pitched as a labor saving device, and even though initial models did not go through a cycle automatically or spin-dry, they did reduce the labor of wash day.

However, washing machines arrived around the same time as a host of other devices, including hot water heaters, irons, and indoor bathrooms. All of these technologies in concert changed users' expectations of "acceptable" hygiene and washing: with so many conveniences, why limit yourself to washing yourself and your clothes once a week?

While individually these devices did save labor, the combination of all of them changed the nature of work in the home. Over time, these devices changed society's expectations about what things would be done, how often, and by whom. Indeed studies of domestic technologies do not show conclusively that work was reduced; more significantly, some suggest that the amount of unpaid work in the home done by women rose dramatically [26].

The washing machine encourages us to take a critical perspective on whether smart home technologies are "labor saving" or whether they, like other devices already at home, merely shift the burden of work. Who will do that work and why?

Other studies show how technologies do not just affect occupant-users, but can become part of broader national debates. Studies of the television and mobile phone show that these devices have influenced how many parents think about "good parenting" [14, 24]. With television, good parenting discussions focus on how much and what kind of programming children may watch. This has, in the United States, led to a broader national debate about the content of television programming. Results of this discussion include a rating scheme for programs, and technologies such as the V-chip.

The mobile phone appears to be taking a similar role in Europe, particularly in countries that have high rates of mobile phone adoption among teenagers and preteens [24]. There, "good parenting" emphasizes two values of mobile phones. First, giving children mobiles helps them learn how to manage bills and money. Second, mobiles allow parents to safely give children increased independence.

As others have noted, smart homes have privacy implications. However, privacy is just one of several social implications of domestic technologies. In this section, we examined two other social aspects of automation in the home: how apparently labor saving devices can actually be labor changing devices, and how technologies influence societal beliefs about good parenting. There are undoubtedly others.

To summarize, there are social consequences that can arise unforeseeably when technology is placed into the home setting. These consequences cannot be reliably predicted from studies of domestic routine, since they alter these routines—and indeed the basic expectations about home life—so drastically. The classic social aspect of computing, privacy, has been explored and addressed by much prior research, from computer-supported cooperative work to ubiquitous computing. But there are other social aspects of domestic computing, as noted here, and their implications can be far reaching. The challenge for us as designers is to be aware of the broader effects of our work, and to realize that even technologies as simple as the washing machine can have broad changes on the dynamics of the home and society itself.

2.6 Challenge Six: Reliability

We can expect that a paramount concern of occupants (if not developers) of smart home technologies is reliability. The range of domestic technologies present in the home today—televisions, telephones, washing machines, microwave ovens—are, by and large, exceedingly reliable, even though these are devices of great complexity. A modern digital television set-top box, for example, contains a number of specialized microprocessors devoted to high-bandwidth decompression, cryptography, rendering, and network communications back to the service provider. And yet, these devices virtually never crash, unlike our desktop computer systems.

Achieving expected levels of reliability, especially when coupled with the ad hoc accretion of devices that may be expected in smart homes, is a great challenge. Dealing with that challenge depends on understanding the reasons that these devices are so much more reliable than "traditional" desktop software systems. Some of these reasons include:

- Differences in development culture
- Differences in technological approaches
- Differences in expectations of the market
- Differences in regulations

First, the development cultures of domestic technologies differ widely from those of desktop, general-purpose computing systems. Embedded systems developers have tended to be much more wary of systems crashes, since it is unwieldy to patch or upgrade a device in the field. A washing machine vendor, for example, would likely fold if it had to recall its products for upgrades as often as traditional software vendors issue patches.

Of course, reliable software systems do exist. These kinds of systems give us insight into how much work it may take to make reliable ubiquitous technologies for the home. Telephone switches illustrate this well; for example, Lucent Technologies

5ESS maintains its reliability goal of 99.9999% (less than 10 seconds of downtime a year) [17]. Meeting this reliability goal means that regular upgrades, such as the ones that provide occupant-users with new services, must be performed while the switch is processing other calls. In other words, this reliability requirement manifests itself within the system architecture. Other parts of the system work on monitoring events that could lead to downtime and either fixing them or reporting them as appropriate [17]. Designing for reliability requires devoting substantial time and resources that will affect the system architecture. Practices such as these must be integrated into the development cultures that will build smart home technology.

The second difference is in the technological approaches taken by domestic technology developers and those in the desktop market. In current connected domestic technologies, the bulk of functionality is placed in the *network*, not in the device itself. In the telephone system, for example, the telephone itself is the least complicated part of the system. And yet it provides access to new functionality available through the network without an upgrade or patch. Digital television systems, likewise, place the bulk of functionality in the network, rather than the client-side device. This is a "utility" approach, in which the client technologies are shielded from upgrades and enhancements in the network, and yet can take advantage of new functionality when available. It is significant to note that embedding intelligence in the network is precisely counter to many of the approaches taken by developers of Internet-based technologies, in which most intelligence resides at the edges of the network. For ubiquitous computing applications, one design challenge will be determining what kind of balance of intelligence to maintain between the edges and the center of the network.

Additionally, the technological approaches taken by designers should account for the need to degrade gracefully. By this we mean that if a component in a richly interconnected system fails it should not bring down the rest of the system. Traditionally, systems have achieved the ability to degrade gracefully through redundancy—data and services are replicated and available on multiple machines. Such an approach may, however, trade off against the goals of simplicity, intelligibility, and ease of administration, which are all requirements for domestic technologies. How to address this tension is a challenge for system designers.

A third difference is simply in the expectations of the various marketplaces. Consumers expect that their appliances will not crash (they have, unfortunately, developed a tolerance for crashes in general purpose computing systems). It is the reliability of so many technologies that has allowed the consumer to actually forget about them as complex technical entities. One hardly thinks of administering the phone or configuring the television. Instead, in large part these technologies blend into the home and become part of the fabric of the home. Crashing phones or televisions would be unwelcome in this setting.

These expectations have been vigorously reinforced by publications and organizations that exist to identify reliable technologies and uncover lemons.[5] Magazines such as *Consumer Reports* in the USA and *Which* in the UK provide information

[5] In the United States cars that do not provide reliably and consequently need repairs early on in their lives are known as lemons. Lemon laws exist to protect the consumer from being sold a "lemon."

about a wide variety of domestic appliances often assessing the reliability as part of their reviews. Organizations such as the Underwriters Laboratories in the USA exist to thoroughly test such technologies before they enter the home.

Fourth, and finally, there are differences in regulation. While the home, as Kidd et al [15] say, is a "free choice environment" for its occupant-users, it is a highly regulated environment for those who provide services into that space. In many Western countries the various utilities that service your home are obligated to deliver a certain level of service, or face regulatory punishments. Insurance companies may demand to see certain levels of safety (such as building upgrades, seismic retrofitting, electrical system changes, and so forth) before they will insure a home. In addition to these *de jure* regulations are *de facto* standards for the home.

All these differences have contributed to services being reliably delivered into the home. Bringing the benefits of ubiquitous computing into such environments may involve creating a development culture that can produce reliable devices consistently, making design choices about how to handle intelligence at the edges of the network robustly, meeting expectations set by other devices, and working toward regulations and standards set by a multitude of agencies. This challenge extends beyond the research community to those who develop, deliver, regulate, and consume these new services.

2.7 Challenge Seven: Inference in the Presence of Ambiguity

Systems in which machine processing is used to control or assist human behavior have a long and less-than-storied track record in the history of computer science. Examples of such systems come from domains as disparate as workflow tools that force users into formal patterns of work [27], and—more recently—Clippit, the Microsoft Office Assistant, which attempts to intuit the actions of a user and offer help.

And yet, much of the literature of ubiquitous computing depicts machine inference of human state and intent as being a crucial factor in the benefits such environments will bring. For example, the literature posits smart meeting rooms that share the notes of the participants [1], and telephone calls that follow their intended recipients throughout a building [25].

Clearly, some of these examples are dependent only on simply detecting and acting on some knowledge of the state of the world from sensor inputs, while other examples are based on a presumption that the system can correctly infer what the user would do him- or herself, if left to his or her own devices.

This begs the question: just how smart does the smart home have to be? How much inference is required for these environments to be successful? What benefits can be achieved with limited inference, or with no inference at all? In the absence of oracular artificial intelligence, how will we design such environments so that their occupants have models about what they can expect their homes to do for them, and how to fix the results of interpretations gone bad?

One constant in published visions of ubiquitous computing is that computing is employed in a physical space to bring functions to users in their everyday work, and that these systems are, to some degree at least, aware of their surroundings, and of their users [25]. The physical world is, of course, what might be termed a “highly analog" environment, presenting a great deal of ambiguity and uncertainty of input—much greater than even the domain of Clippit.

Intelligence in such a world can take a number of forms, some of which make greater assumptions than others. Some of the more obvious of these include:

- The environment can interpret the meaning of sensor data to reflect some state of the world. For example, the system might assume that I am in a room because my active badge is in a room.
- The environment can infer that some state exists by aggregating a number of other factors. For example, if a number of people are gathered together in a meeting room, the system might assume that a meeting is taking place.
- The environment may attempt to infer my intent from its view of the state of the world. For example, the system might assume that because I am in a meeting, I might want to share my meeting notes with others in the meeting.
- Finally, the system may preemptively act on assumptions of intent. For example, if the system assumes I may want to share my meeting notes, it may go ahead and make those available to other meeting participants (or ask me if it should do so).

All of these modes of intelligence can be found represented in the literature of ubiquitous computing (see [6], for a similar categorization). And all are subject to error, of varying degrees and types.

For example, the simple sensing case may report that I am present in a room when, instead, I have simply left my active badge on the desk. These are what might be called "phenomenological" problems—do sensors reflect reality or merely the state of the sensors—and can, in all likelihood, be largely overcome by more and better sensor technology (although perhaps at a cost of privacy and user control). And—perhaps more importantly—the cost of incorrect inferences is low if the system does little with the inferred information.

More dramatic problems become apparent as uncertain inferences and decisions are compounded. Most troubling is the attempt at inferring some internal human intent and then, perhaps, taking action on it, especially when such an inference is based on layers of ambiguous interpretation and input, or requires a level of intelligence that even humans would find difficult.

Our challenge, then, is to discern what functions of the smart home are possible with limited inference, which are possible only through inference, and which require an oracle. The first category comprises good candidates for implementation, since limited machine interpretation means that there is limited possibility of error. The third category, systems that require omniscient understanding of human intent in order to function well, are perhaps better abandoned.

The middle category, we feel, is the most interesting, and presents important problems in design and technology. Systems that rely on inference will never be right all of the time, and thus users will necessarily have to have models of how the system arrives at its conclusions. These models must not only concern themselves with the actual rules of inference ("when people gather in the living room, display the television schedule"), but also the capabilities of the system's sensors ("how does the system know I'm in the living room in the first place?").

Users must know what to expect from their homes in the same way that, say, a user knows that dropping temperature outside will cause the thermostat to turn on the heating [4]. Such predictability depends on:

- The system's expected behavior in the face of this condition is known.
- The system's facilities for detecting or inferring this condition are known.
- Provision is made for the user to override the system's behavior.

Achieving these three conditions is more complicated when the inferences made by the system are more complex, and when even basic sensing is unreliable or open to interpretation.

The challenge for smart home designers is to create systems that ensure that users understand the pragmatics of sensors, interpretation, and machine action as well as they understand the pragmatics of devices in their homes now. From a technical perspective, the challenge of developers is to ensure that ambiguity is not hidden from the parts of the system (or the users) that need access to it, and to ensure that inference—when performed at all—is done in a way that is predictable, intelligible, and recoverable.

3 Discussion

In this paper, we have presented seven challenges for ubiquitous computing research in the home setting. Although we have divided these challenges for the purposes of discussion, there are clearly interesting connections and overlap among them.

First, there are a number of problems that arise that are unique to the smart home setting itself. Existing houses were not designed to be smart, and office technologies are often intended to be just that—technologies for the office rather than the home. The realities of the home setting, coupled with the fact that adoption of home technologies is likely to be incremental and disjoint for the foreseeable future, give rise to our first and fourth challenges.

Second, we believe that there are a host of technical, implementation, and systems design issues that are often underestimated. Further (and perhaps more importantly) the *tradeoffs* among these issues are not well understood in the home setting. The stringent requirements for reliability, the fact that there will be no formal systems administrator (nor will most homeowners be likely to want to undertake such a role), and the desire for interoperability all trade off against each other. For example, it is relatively easy to make an easily administered and reliable device if that device never needs to communicate with or use any other services in the home. We believe that finding the right balance between these requirements is crucial.

Third, and as always, the social impact of new technologies is hard to predict. The home setting is not novel in this respect, although the social dynamics and relationships within the home make it perhaps a more volatile setting than the office or other public spaces.

Finally, we believe that an overarching philosophical question that should be addressed by designers of smart home technology is, simply, how smart does the smart home have to be to provide utility to its occupant-owners? To some degree, this question permeates all of the others we have raised, since it is precisely the "smartness" of the smart home that makes it disruptive to the domestic order, gives rise to the architectural and implementation tradeoffs mentioned above, and makes the social adoption of this technology so unpredictable.

We believe that the chief challenge that will be faced by the designers (and, potentially, the occupants) of the smart home is balancing the desire for innovative technological capabilities with the desire for a domestic lifestyle that is easy, calming, and—at lease in terms of technology—predictable.

4 Conclusions

This paper has presented seven challenges that we believe must be successfully addressed for the smart home to be a viable place to live. These challenges span across technical, social, and pragmatic domains.

Our hope in raising these issues is to create a conversation among researchers in the varied disciplines that make up ubiquitous computing. In particular, we hope to raise awareness of the existing literature on the adoption, use, and history of domestic technologies, as well as the use of situated studies, and the benefits that these can bring to bear on the design and evaluation of technologies for the home.

Acknowledgements

We would like to thank our PARC colleagues who challenge us to think more clearly about the challenges confronting ubiquitous computing. Special thanks to Diana Smetters for the wireless speakers example.

References

1. Abowd, G.D.: Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment. IBM Systems Journal, 38 (4). (1999) 508-530
2. Abowd, G.D.,Mynatt, E.D.: Charting Past, Present, and Future Research in Ubiquitous Computing. ACM Transactions on Computer-Human Interaction, 7 (1). (2000) 29-58
3. Arbaugh, W.A., Shankar, N.,Wan, Y.C.J.: Your 802.11b Wireless Network has No Clothes. University of Maryland, Department of Computer Science (Technical Report). (2001)
4. Bellotti, V.,Edwards, W.K.: Intelligibility and Accountability: Human Considerations in Context Aware Systems. to appear in Human Computer Interaction, 16. (2001)
5. Buderi, R.: Computing Goes Everywhere. Technology Review, Jan/Feb. (2001) 53-59
6. Dey, A.K.: Understanding and Using Context. to appear in Personal and Ubiquitous Computing, 5 (1). (2001)
7. Edwards, W.K., Newman, M.,Sedivy, J.Z.: The Case for Recombinant Networking. Xerox PARC Technical Report (2001)

8. Fischer, C.S.: America Calling: A Social History of the Telephone to 1940. University of California Press, Berkeley, CA (1992)
9. Fox, A., Johanson, B., Hanrahan, P.,Winograd, T.: Integrating Information Appliances into an Interactive Workspace. IEEE Computer Graphics & Applications, 20 (3). (2000) 54-65
10. Georgia Institute of Technology.: Aware Home Research Initiative. http://www.cc.gatech.edu/fce/ahri/.
11. Grinter, R.E.,Eldridge, M.: y do tngrs luv 2 txt msg? In Proceedings of the Ninth European Conference on Computer-Supported Cooperative Work ECSCW '01. Bonn, Germany, Dordrecht, Netherlands: Kluwer Academic Publishers, (2001)
12. Hales, L.: Blobs, Pods and People. The Wall Street Journal, Sunday March 25. (2001) W34
13. Harvey, D.: Broadband Security: Who's Responsible? Home Office Computing Magazine, August. (2000)
14. Hughes, J., O'Brien, J.,Rodden, T.: Understanding Technology in Domestic Environments: Lessons for Cooperative Buildings. In Proceedings of the First International Workshop on Cooperative Buildings (CoBuild'98). Darmstadt, Germany, Heidelberg, Germany: Springer-Verlag, (1998), 248-261
15. Kidd, C., Orr, R.J., Abowd, G.D., Atkeson, C.G., Essa, I.A., MacIntyre, B., Mynatt, E.D., Starner, T.E.,Newstetter, W.: The Aware Home: A Living Laboratory for Ubiquitous Computing Research. In Proceedings of the Second International Workshop on Cooperative Buildings (CoBuild'99). Pittsburgh, PA, Heidelberg, Germany: Springer-Verlag, (1999)
16. Kindberg, T.,Barton, J.: A Web-Based Nomadic Computing System. HP Labs Technical Report HPL-2000-110, http://cooltown.hp.com/papers/nomadic/nomadic.htm. (2000)
17. Lucent Technologies.: Press Announcement: 5ESS Reliability. http://www.lucent.com/press/0699/990602.nsb.html.
18. Markoff, J.: Flaw in Popular Wireless Standard. New York TImes, http://www.nytimes.com/2001/04/03/business/03FLAW.html.
19. Norman, D.A.: Invisible Computer: Why Good Products Can Fail, the Personal Computer is so Complex and Information Appliances are the Solution. MIT Press, Cambridge, MA (1999)
20. Nye, D.E.: Electrifying America: Social Meanings of a New Technology. MIT Press, Cambridge, MA (1990)
21. OSGi: Open Services Gateway Initiative. http://www.osgi.org/.
22. Palen, L., Salzman, M.,Youngs, E.: Going Wireless: Behavior and Practice of New Mobile Phone Users. In Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW 2000). Philadelphia, PA, New York, N.Y.: ACM Press, (2000), 201-210
23. Race, T.: Adding Firewalls to Home PCs. The New York Times, March (12). (2001)
24. Rautiainen, P.,Kasesniemi, E.-L.: Mobile communication of children and teenagers: case Finland 1997-2000. In Proceedings of the Workahop on "The social consequences of mobile telephony: the proceedings from a seminar about society, mobile telephony and children". Oslo, Norway, (2000)

25. Schilit, B.N., Adams, N.I.,Want, R.: Context-Aware Computing Applications. In Proceedings of the Workshop on Mobile Computing Systems and Applications. Santa Cruz, CA, IEEE Computer Society, (1994), 85-90
26. Schwartz Cowan, R.: More Work for Mother: The Ironies of Household Technology from the Open Hearth to the Microwave. Basic Books, Inc., New York, NY (1983)
27. Suchman, L.: Speech Acts and Voices: A Response to Winograd et al. Computer Supported Cooperative Work: An International Journal, 3 (1). (1995) 85-95
28. Weiser, M.: The Computer for the Twenty-First Century. Scientific American. (1991) 94-104
29. Weiser, M.,Brown, J.S.: Designing Calm Technology. PowerGrid Journal, 1 (1). (1996) http://powergrid.electriciti.com/1.01