Home Networking and HCI: What Hath God Wrought?

Erika Shehan and W. Keith Edwards
School of Interactive Computing / GVU Center
Georgia Institute of Technology
85 5th Street NW
Atlanta, GA 30308
{erika, keith}@cc.gatech.edu

ABSTRACT
For much of the industrialized world, network connectivity in the home is commonplace. Despite the large number of networked homes, even the most technically savvy people can have difficulties with home network installation and maintenance. We contend that these problems will not disappear over time as the networking industry matures, but rather are due to structural usability flaws inherent in the design of existing network infrastructure, devices, and protocols. The HCI community can offer a unique perspective to overcoming the challenges associated with home networking. This paper discusses why home networking is difficult, based on analysis of historical, social, and technical factors. It explores how the designs of existing home networking technologies have implications for usability, and examines a range of models for addressing these usability challenges. The paper concludes with a discussion of how these models may impact future research efforts in both HCI and networking.

Author Keywords
Computer networking, home networking, home computing, infrastructure, troubleshooting, ubiquitous computing

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous. K8.3. Computing Milieux - Personal Computing: Management/Maintenance

INTRODUCTION
In 1844, Samuel Morse transmitted the first message over the newly invented telegraph. The message read "*What hath God wrought?*" At the time of that message, Morse was unaware of the social, economic, and technical factors that would ultimately shape adoption of his new technology. Telegraphy would evolve in unexpected directions; it would also give rise to entirely new professional classes and social practices that co-evolved alongside the new technology [32]. Today, a different set of social, economic, and technical factors are shaping the adoption and practices surrounding another relatively new technology: computer networking. Much like Morse's telegraph in the 19th century, these forces have pushed networking into places unforeseen by its original inventors.

One of those unexpected places is the home. As of late 2005, 43.2 million US households had a broadband connection, and nearly 30 million of these households also had a home network [36]. While the numbers in other nations may vary, for much of the industrialized world home Internet connectivity is a common feature. In many of these homes, networking is not simply one computer directly connected to a broadband connection, but rather consists of complex networks built *within* the home [36]. Given such a large number of households with home networks, one may think that the problem of networking in the home is solved, or even that there is no problem at all. Data from consumer and industry marketing organizations, however, point to a different story. With return rates of 20-30%, home networking gear is currently *the most returned item* at "big box" electronics stores [30]. As recently as 2006, roughly a quarter of wireless access points purchased by consumers were returned [23]. Much of the reason for these returns is not due to technical failings of the devices themselves, but rather due to user experience failings. For example, the installation and configuration is so difficult for many users that they cannot correctly get these devices to work at all. In fact, consumers cite technical complexity as the largest barrier to home networking [20]. These examples point to deep problems in the usability of networking for home users.

In this paper, we explore the usability issues in home networking, arguing that many of these are inherent in the current technical and economic development of home networking. Resolving these problems will require the involvement of the HCI community in shaping the usability of current and future networking technologies. We provide glimpses of several possible futures of home networking from a perspective that comes primarily from HCI, but is informed by networking. We conclude with a discussion of possible future research directions in this area.

CHI 2007, April 28–May 3, 2007, San Jose, California, USA.

UNPACKING THE TRAVAILS OF NETWORKING

In this section we examine the impacts of networking on the user experience, and explore how historical and technological features in the current home networking stack make the installation, setup, and maintenance of home networks so vexing for householders.

First of all, networking is fundamentally an *infrastructure* technology. Although people may *use* networked applications such as email or instant messaging, they generally do not *see* the network itself (apart from physical devices such as routers or cables) during normal, functional use. Networking, like all other infrastructure, is, in the words of Star, "invisible, part of the background for other kinds of work" [33]. Despite this invisibility during normal operation, networks have a direct, tangible impact on how home users interact with the systems built atop them. This impact is most obvious when networked technology malfunctions. In these situations, users must understand and interact with networking in order to correct the problem. Networking also becomes visible when users configure and install new devices, change the behavior or parameters of the network or its components, or remove a device from the network. In all of these cases, the previously invisible infrastructure is *foregrounded*, and users perceive and experience the infrastructure directly [5].

Second, the multiplicity and complexity of the devices on the network means that creating and maintaining a robust home network is cumbersome, error-prone, and fragile. Seemingly simple changes to device settings or network configurations can lead to network failure and hours of troubleshooting. Common causes of network trouble may include physical component failure inside the home (e.g., a physically broken or unplugged device), incorrect configurations on a device (e.g., incorrect IP address), interference from other electronics (e.g., cordless telephones), or even inherent incompatibility between devices on a network (e.g., due to different versions of specs such as UPnP, or different codecs used in media encoding). Failures may also originate from points outside of the home (e.g., trouble at the user's service provider).

Third, the current state of home networking requires that householders deal with a number of players outside of the home, including the service providers who deliver internet connectivity, as well as vendors of hardware, software, and services. Most software and hardware used in PCs and home networking equipment are not built with expectations of 100% reliability; mysterious program crashes and malfunctioning devices are the norm rather than the exception in home computing. The sometimes incompatible interplay between these different entities creates a confusing situation for householders, who often don't know where to turn for help when solving problems [28]. These difficulties prompt a range of coping strategies. Some people rely on social networks of technically sophisticated friends and neighbors for troubleshooting advice [19]. (In fact, the book *Networking for Dummies* recommends that users "start with friends, relatives, or neighbors who are computer geeks" to get help with networking [22]). Other people, frustrated with their inability to maintain the networks they have deployed, turn to paid external professional help to repair the network. The price some people are willing to pay for technical help with their digital devices tells a story of difficulty with computing. A single visit from the Geek Squad (an at-home tech support service in the US) to remove spyware and viruses can cost as much as purchasing a new computer [11]. Finally, in the worst case, a number of people decide that a home network is simply not worth the trouble and give up completely. At this point, one might wonder how, exactly, did home networking get to this state? In the next sections we first examine the causes for this difficulty, which stem from the architectural and historical considerations that drove the design of networking technology. We then look at the effects of these considerations on user experience.

Causes: A Brief History of Networking

Home networking technology uses the same packet-switched protocols of the larger Internet. These protocols were originally designed for use by technology researchers and the US military in the midst of the cold war. (Indeed, the first Internet technologies were designed to survive in the face of a nuclear weapon attack [14]). Important design goals of these early networked technologies included scalability (the ability to add more machines), extensibility (the ability to develop new sorts of applications upon a common infrastructure), and throughput (the ability to transfer data at adequate speeds). These design choices are reflected in current home networking technology, and are undoubtedly positive features even in small home networks.

A number of networking researchers, however, have begun to question the appropriateness of many of the design assumptions of these protocols, in light of how network usage has changed in the last 40 years (see, for example [4]). For example, in the early days of the Internet, all machines on (and people using) the network were generally considered trustworthy; thus no low-level authentication or access control mechanisms were built into the core Internet protocols. Likewise, the basic structures of IP and TCP/IP make no guarantees about quality of service, which can present problems for entertainment applications for streaming audio and video data. [4]. Moreover, because of the guaranteed technical expertise of the people using the predecessors to the Internet, ease-of-use at the network endpoints (actual host computers) was not a primary consideration when designing these protocols [1].

Further, one of the major design goals of this technology was that "specific application level functions usually cannot and preferably should not, be built into the lower levels of the system—the core of the network" [4]. This design choice has direct implications on the usability of current home networking technology. We strongly acknowledge that this design choice has a number of important benefits. Most importantly, it greatly simplifies the design of the core of the network: by limiting the capabilities built into the network itself, the network core can stay relatively simple and fixed, requiring few upgrades while still supporting unlimited extensibility at the edges of the network. This design choice has proved its value over and over again, as it has allowed a range of applications (such as email and the web) to arise without requiring any changes to the core Internet routing infrastructure or protocols; adding new functionality only requires agreement at the endpoints (such as the SMTP and HTTP protocols, in the case of email and the web, respectively), not in the network core [4].

A negative consequence of this design choice, however, is that functionality is pushed out of the core of the network to the network edges—in other words, to the components and devices that are installed in users' homes. Because the network design requires that client devices must be correctly configured in order to communicate via the relatively "dumb" network core, someone (or something) *must* do this configuration. Further, where there is the possibility of configuration, there is the possibility of mis-configuration, which in the case of the Internet often prevents client devices from communicating at all. In the Internet model, client devices are largely stateful (since they must maintain their configuration information), complex (since they must be capable of dealing with an open-ended set of application-layer protocols), and managed (since the device must provide capabilities for some one or some thing to configure it correctly for the network). Rather than reducing the costs of operating a network, these design choices push the cost and complexity of networking into the hands of the householder.

Note that these usability properties are not determined simply because networking is a communication technology. For example, the public circuit-switched telephone network has a radically different infrastructure in which intelligence is shifted into the network core and away from the edges. For traditional wired phones, this arrangement provides a radically different user experience. Phones are not stateful devices; they do not "know" their phone numbers, but rather automatically acquire them from the network. Phones are generally simple, single purpose devices that only must understand one simple protocol (and yet, once connected, can place calls to virtually every other number in the world including to devices such as mobile phones that may not have even existed when the wired phone was built). Telephones are not managed (since generally there is nothing that needs to be managed). The user experience of such devices is that a user simply plugs them in and they work [8]. These examples show how architectural and protocol design decisions can deeply affect usability.

Effects: Challenges of the Home Environment

When Internet-style networking protocols are put into the home environment, a number of unforeseen problems appear. Some of the biggest problems of networking in the home environment are related to *statefulness*. Devices must be configured with detailed information that is often difficult to remember (machine name, IP address, components installed on the device, details about patching, etc). Rarely do people keep documentation of the structure of their networks; a previous study shows that even in the same household, occupants may have very different conceptualizations of the *same* network [13].

Additionally, people often do not plan or manage for network growth. Devices are often added in a haphazard, piecemeal manner. Moreover, it is difficult to tell when and where changes to settings have been made. A curious child who makes a few haphazard keyboard presses can wreak havoc upon the settings of a network. Overall, there are few, if any, tools that allow for structured management of the home network. The open-ended nature of home networking protocols has also led to highly dynamic and heterogeneous home networks. The range of possible devices—and possible ways to configure and connect those devices—means that every network likely looks different. Further, home networks are highly dynamic; devices connect to and disconnect from home networks far more often than in large enterprise environments [38].

Since always-on broadband connections have come into the home, *security management* has become a significant problem. The trust assumptions of the original Internet protocols cannot be assumed today [2], as evidenced in the high rate of network security exploits. Automatic patching can assist in overcoming this problem, but it also can break working configurations, and in some cases, machines may be infected more quickly than patches can be installed. Despite the widespread availability of add-on antivirus programs for end user devices, malware such as viruses and spyware are still a nuisance. Network firewalls can offer protection against security risks, but they also are another opportunity for misconfiguration, and require knowledge of networking technology in order to use effectively.

All of these problems are exacerbated by the fact that there are very few tools available to help users who may not have significant technical knowledge or motivation fix their networking problems. Common tools (e.g. ping, nslookup, traceroute) are inaccessible to the average user, and troubleshooting techniques are often not well-understood by people who are not technicians by profession. Often rebooting everything on the network is the only understandable strategy available to householders.

Division of labor within the household can also magnify networking difficulties. A previous study shows that knowledge about the network is often inequitably distributed between members of a household [13]. This inequity can lead to problems when knowledgeable members of the household aren't available; in one household studied by [13], an occupant went without networking for a week because the only person who understood the network's configuration was out of town.

ALTERNATIVE APPROACHES TO HOME NETWORKING

In the previous sections we have explored how design choices made in network architectures have user experience implications. Given the inherent architectural properties of the Internet, are there ways to make networking less painful for householders? Are there approaches that might maintain the overall goals and aims of the Internet while offering radically improved usability?

In the following sections we explore a number of possible futures for networking technology in the home. Each of these futures embodies a set of social, historical, economic, and technical outcomes that not only provide different usability properties, but also have implications for networking and HCI research directions, the degree of agency or control which users may have over their networks, and how economic benefits are distributed. These outcomes are not necessarily mutually exclusive; in fact, there is significant overlap among several of them. Each, however, represents a distinct model for how networking in the home may evolve; each has also been proposed (perhaps implicitly) by research and industry, although not always under the names we give them here. The possible futures we examine are: (1) The Fresh Start Model, (2) The Bandage Model, (3) The Gateway Model, (4) The Outsource Model, (5) The Utility Model, and (6) The Ubiquity Model. For each of these models we describe the outcome, the opportunities it represents, and the factors that mitigate for or against it.

The Fresh Start Model

We have argued that many design choices in the packet switched network architecture have led to poor usability of home networking. One might reasonably ask whether the only way out of the problem is to abandon the architecture altogether, and consider one that can provide a better user experience "built in." Since the uses of Internet-style networked technology have far outgrown their original purpose of military, research, and government use, what if we were to simply take everything we have learned over the course of developing networked technologies over the past 40+ years, and rebuild networking infrastructure so that it is easier and more appropriate for use in homes of the present—and the future?

This *Fresh Start model* offers opportunities to rethink networking protocols from the ground up, potentially offering "deeper" usability than is currently possible—ensuring that networks are easier for people with little technical expertise to use, that they are flexible for a range of new devices, and do not make assumptions that all devices on the network are trustworthy. A number of research efforts (such as NSF's GENI [26]) are exploring this approach, focusing on attributes such as ease of configuration and security. While this approach is a worthy and necessary direction for research, it is, however, unlikely to be achievable in any short time span (or perhaps ever). Any infrastructure has inertia: once deployed it can be difficult to change significantly, because of the need for backwards compatibility. This effect has variously been called *soft determinism* by David Nye [27] and *path dependence* by economist Nathan Rosenberg [29]:

> "Early pressures and decisions can 'lock-in' a certain technology or market structure for a far longer period of time than may be socially optimal…it is often very difficult and costly to reverse technology decisions once they have been made.... The ability to evaluate current technological options is complicated by an inability to predict the future evolution of each technology before its adoption." (p. 205)

We have already seen the great weight of path dependence in the Internet, where the network effects of scale work against rapid change. Despite push from technology and governmental interests, the core Internet protocol is still at Version 4 [17] (which was first introduced in 1981), despite the widespread availability of Version 6 (the most recent) on client devices [18]. Likewise, HTTP—the protocol that must be agreed upon by web servers and browsers—is only on Version 1.1, and even that is not completely widespread [16]. It is arguable whether a Version 1.2 will ever exist.

Because of path dependence, replacing the existing Internet with something incompatible with it would require en masse (and perhaps simultaneous) adoption of the new technology, not only in homes but also in enterprises, governments, and the network core itself. A change of this scale presents a significant coordination problem in the commercial sector, because it would require hardware and software manufacturers as well as ISPs to jointly agree to a particular type of upgrade; this task would be difficult, if not impossible [2]. The last time this sort of upgrade happened was in 1983, when all Internet-connected machines switched to TCP/IP simultaneously in an "Internet flag day." Although only a few hundred machines were involved with the transition, it took several years to plan for the event [14]. It seems unlikely, if not impossible, that this sort of transition could happen today, given that there are millions of machines that would have to simultaneously upgrade.

Of course, it may be possible to upgrade the Internet-style technology in a way compatible with existing tools, which might not necessitate a large-scale simultaneous upgrade. While such a technological solution is perhaps imaginable, it would—by definition—have to maintain features of the current architecture, many of which are detrimental to usability. How many of the usability problems such an upgrade could potentially solve remains an open question.

The Bandage Model

There are, of course, smaller upgrades to network technology that may not require replacement of existing technology, which we refer to as the *Bandage model*. Such approaches may work around usability flaws through clever UI design and extensions to existing technology, keeping the underlying infrastructure while building new interfaces and tools that help users understand, repair, configure, and use their networks. While not providing the top-to-bottom approach to usability that the Fresh Start model might allow, such tools could still potentially lower the barrier to “self-service” in the home, empowering householders to take care of complex networking chores themselves without having to resort to outside assistance such as a knowledgeable friend, telephone support line, or paid troubleshooters. Alternatively, these tools could support those external troubleshooters in doing their jobs more easily and cheaply.

Already there have been several commercial and research attempts at providing better tools for network setup and maintenance that work within the existing TCP/IP architecture. Perhaps most widely known are technologies such as the Dynamic Host Configuration Protocol (DHCP) and various discovery protocols [9]. These technologies take care of automatically assigning device state, and providing administration-free detection of peer devices, respectively. These “alongside” technologies provide an improved user experience. Further, they work with the existing Internet infrastructure. For example DHCP-enabled devices can coexist on a network with non-DHCP-enabled devices. There is no need for wholesale buy-in of the technology in order for benefit to accrue.

Research and commercial work has also been done on interface-layer improvements for tasks such as network setup. Most of these introduce a centralized component into the network that has responsibility for configuring clients and maintaining their state; in effect, clients delegate their setup tasks to the centralized component, representing a small-scale shift back to the “intelligence in the network” model. These systems include the Linksys EasyLink Advisor [23] and PARC Network-in-a-Box [3]. Others provide new interaction techniques for exchanging the configuration information necessary to work within the existing network architecture. These systems, for example, may send information necessary for a laptop to join a wireless network via infrared rather than through manually-entered hexadecimal keys. These systems include Microsoft Windows Connect Now [24] and Linksys Secure Easy Setup [23]. Yet another set of systems aim to ease network troubleshooting within the home. These include Network Magic [25] and the Georgia Tech network visualization project [31]. These systems, built on top of the protocols already provided by the network, show interactive visual representations of devices and network interconnections to help users understand their networks and resolve problems.

This approach to improving the usability of home networking offers promise. Home networking setup and maintenance is frequently done by those with substantial “do it yourself” network skills [13], or those with close connections to people with those skills [22]. Empowering relatively unskilled and unmotivated users to become effective “do-it-yourselfers” in the home space seems a worthy goal; by allowing more people to fix common problems more easily, we can perhaps remove many of the barriers to maintenance. Note that these tools need not be only for the householders themselves. Indeed, this approach offers new opportunities for businesses, which may be able to extend and improve their technical support tools to (1) allow users to deal with small problems themselves (which is especially important for low-margin consumer electronics producers, for whom a single support call may negate the profits derived from the sale of a product), (2) provide more precise information to remote troubleshooters, or (3) have a collaborative, shared view of a network to assist with troubleshooting.

While this approach offers substantial opportunity, it has significant limitations. Fundamentally, tools such as those described here must be compatible with the existing home networking infrastructure. Although these tools can hide many of the more painful aspects of home networking, they are often only skin deep. Many network analysis tools, for example, depend on being able to collect data reliably from multiple points on the network. A network that is partitioned or otherwise broken beyond the ability to communicate may make it impossible for most tools to deliver reliable troubleshooting guidance. Furthermore, visualization tools can only visualize data that the network can actually provide. If the underlying home network infrastructure does not provide access to the details necessary for problem solving, a tool built on top of this infrastructure cannot create it. Achieving better, deeper usability may require more fundamental changes to the underlying networking technology.

The Gateway Model

There is a feature of the Internet architecture that may lend itself to a better user experience for home networking: isolation of sub-networks from the greater Internet. This feature means that individual networks (such as one in a home) can run on entirely different protocols, with entirely different networking infrastructure than the Internet at large. A single gateway device can translate between the home network and the Internet, meaning that only the home networking infrastructure, rather than the entire infrastructure of the Internet, needs changing. This *Gateway model* has the possibility to provide the same sort of “deeper” usability as the Fresh Start model, although without the need for substantial change at the network core. Since the home network is likely to be a different sort of networking environment than the Internet as a whole, different design choices can be made. For example, Internet-scale scalability can be sacrificed if necessary to obtain better usability. Likewise, it may be possible to embed stronger trust assumptions in the devices on the home network.

By definition, this model requires a smaller set of changes than the Fresh Start model—only householders would have to adopt the new technology, not other players in the network ecosystem. Further, this model avoids the chief obstacle to the Fresh Start model, in that adoption does not require coordinated change across the entire network:

individual homes can adopt the revised technology in a piecemeal fashion. Finally, although this model relies on a single point of failure (if the gateway breaks, no communication between devices and the outside world can occur), this situation is no worse than in current networks if the router or cable modem dies.

Of course, this model does come at a cost to householders in the form of broken compatibility with existing devices. In order to work with the new technology, householders would have to upgrade their devices. These changes can be annoying, but the promise of more usable, reliable technology may be enough to convince some people to switch. The piecemeal adoption afforded by this model also means that different households can exist at different points on the adoption curve simultaneously.

The Outsource Model

One obvious possible future for home networking—and one that has substantial historical precedent—is the *Outsource model*, in which householders turn to external professional help to install, maintain, and manage their networks. A number of other domestic technologies have followed this path, as the types of "production work" and "consumption work" done in the home have changed [7]. As noted earlier, this trend has already begun with services such as GeekSquad in the United States. If we believe that networking technologies will also follow this well-worn path, then this suggests a range of areas for both networking and HCI work that mirror trends of use of past technologies. For example, we suggest that developing new technologies that improve the ability of outsiders to troubleshoot and repair the network may be fruitful, in the same way that onboard diagnostics in modern cars improve the ability of car repair technicians to diagnose problems in automobiles.

Despite the historical precedent for outsourced maintenance, technical characteristics of networking—as well as social factors in the home—mean that home networking may not be as amenable to an outsourcing model as appliances and cars are. Some of these differences, of course, arise simply from the newness of networking in the home: if the Outsource model becomes prevalent, and thus market forces generate more repair services, professional infrastructure (such as ratings and accreditation services) will likely arise in due time. Other differences, however, are more fundamental in nature, for reasons including heterogeneity, lack of technical support for diagnosis and repair, and privacy.

Overall, appliances and cars are more limited in variety than home networks are. Additionally mechanics have service bulletins and reference materials from the manufacturers or third parties that that show car layouts and describe common repairs. Technicians repairing a computer network may have manuals for each individual component on the network, but they do not have information about configurations specified by householders, or about how all of the devices are connected together. Furthermore, the auto repair industry does not have to deal with cars that are user created and configured (by often inexperienced designers), cobbled together with the technological equivalents of bubblegum and shoestring.

Given that there is little diagnostic information to help the computer technician, network repair is a much more difficult problem. One solution to this issue is offering professional, standardized networking setups. Another solution would be to have shared views of networks that can be used for remote problem diagnosis and repair. Outsourced maintenance, however, may present privacy issues. Unlike repairing a hot water heater or car, computer and network repair requires providing a technician with access to deeply personal information stored on personal digital devices. Will people be comfortable with disclosing this much information to strangers?

The existing home networking repair services demonstrate the economics of this model, as paid outsourcing is already creating a new market. Outsourced maintenance may be a reasonable solution for householders who do not care about networking and would prefer to pay someone else to think about networks for them. However, not everyone who uses a "do it yourself" approach to home networking does so because he or she enjoys network maintenance. Some people simply have no other alternative because they cannot afford the price of outsourced maintenance.

The Utility Model

Some have proposed [2,6] the future of networking by describing its user experience: "it should just work." Here, we characterize this vision as the *Utility model*, drawing on similar experiences of utilities in the home such as electricity, natural gas, and the telephone.[1] These utilities have a number of desirable usability qualities: the experience of electricity, for example, is that users plug devices into sockets and they *just work*. Even complex distributed communications systems—such as the telephone and cable networks—largely have this property. Could we achieve such a level of usability with home networks?

Again, there are fundamental differences between these services and home networking that challenge whether a utility model is realistic. These differences center on core assumptions of the Internet-style architecture, the wealth of customization options supported by rich client devices, and the visibility of the current networking infrastructure in contrast to current utilities. Although utilities such as electricity, natural gas, and the landline telephone "just work" in the home, these utilities differ greatly from computer networking. Electricity, natural gas, and the landline telephone system are highly regulated, closed

[1] We differentiate this model from that of current ISPs, even though these are also described as utilities. Because of the way the network is structured, the responsibilities of current ISPs end at the network router—the services *inside* the home are not within the ISP's purview.

systems. All of the “intelligence” of these systems is in the network itself, not in the end user devices. This closed, regulated, network-intelligent design is counter to that of the design and culture of the Internet [4]. For example, in the early days of the AT&T telephone service monopoly in the US, customers were not allowed to attach third party equipment to telephone lines due to concerns that these devices could have detrimental effects on the entire telephone system [15]. This tight regulatory control is rare in the world of computer networking (with the exception of ISP terms-of-service agreements).

Data formats present another challenge for networking to be like a utility. Consider the case of natural gas: gas doesn't come in a variety of confusing human-created formats, nor is it something that can be configured. There aren't different gas line protocols, and connection between the gas infrastructure and the home is far less complicated than the typical cable/DSL hookup. Besides the dangers of gas leaks, users need not worry about security of their gas hookups; given the highly regulated, closed nature of the gas system, householders do not need to be concerned with the possibility of malicious parties inserting “bad gas” into their lines that would break their stoves or heaters.

Additionally, utilities offer householders few opportunities for customization and configuration. When one connects an appliance to a gas line, the item does not need to be configured or customized. Additionally, there are no quirks related to standards or manufacturer specific implementations of a standard. Moreover, utilities such as gas or electricity may power an appliance, but have little impact on its overall user experience, aside from the possibility of a fuse being blown, or a pilot light being burnt out. In contrast, networking problems can be much more subtle than “the appliance won’t turn on” or “the pilot light is out.” But neither the networking nor the HCI communities have strong approaches for dealing with such usability problems that are a result of issues deeper than the interface level [10].

Finally, compared to utilities, people are far more likely to tinker with and customize the settings of their networks. Unlike with gas or electricity, where mistakes may lead to property damage or death, computer networks do not generate the same level of cautiousness. Furthermore, in order for networks to be *useful* (as opposed to merely *usable*) they must accommodate the social structures and patterns of home life [8]. This implies that networks will likely always require more opportunities for customization and personalization than utilities such as gas or electricity.

The Ubiquity Model

Similar to—but subtly different from—the Utility model is the *Ubiquity model*. In this model, the “home” network does not exist at all. Rather, wireless networking is provided as a ubiquitous service by either a commercial or community-sponsored entity. In this model, although householders would have to manage the connectivity of their devices to the network, they would not have to concern themselves with tending to details of the network infrastructure itself. The technological foundations for the Ubiquity model are usually framed in terms of municipal-scale WiFi, 3G data services, or—more recently—WiMax-based systems that can provide high-bandwidth coverage over a metropolitan area (see, for example, [39]).

This model has a number of benefits. First, a large-area wireless network brings the potential for easy mobility. Second, by doing away with the need to wire the “last mile” connections to the home, this model may lower barriers to access for underserved communities that may not have infrastructure for conventional broadband. Finally, this model also opens up new opportunities for collaborative applications that leverage the fact that an entire community, not just a single household, shares a network.

On the other hand, many components in the networked home—printers or large networked appliances, for example—do not benefit from the mobility advantages of this model. More subtly, the Ubiquity model does away with aspects of the (already weak) security protections provided by current home network architectures. In the majority of current home networks, there is a security boundary at the edge of the home network provided by a firewall or router. This boundary ensures that devices on the home network are only accessible from *within* that network. With this sort of boundary, devices inside the home can generally have weaker defenses, as they are assumed (by virtue of the fact that they are on the same physical network) to be trustworthy. The Ubiquity model, however, does away with the basic assumption that devices on the same network share a trust relationship. In a world in which an entire metropolitan area is on the *same* network as individual home devices, householders need new mechanisms to “add back the walls” of the home network. These mechanisms would allow them to define trust relationships that are not based on network proximity, and provide defenses for devices (e.g., printers) that were not previously exposed outside of the home.

Finally, in many countries, the means by which the Ubiquity model may play out will likely be determined by political processes rather than technical ones. In the United States, for example, there is no consensus on whether networking should be a public good or a profitable, commercial entity. Whose interests will speak louder—those who may otherwise be denied access or those who are running a business?

LOOKING FORWARD: FUTURE RESEARCH AGENDAS

The previous sections have examined possible futures for networking technology in the home, based on historical precedent, social properties of the home, and fundamental attributes of current networking architecture. All of these models suggest directions for future HCI and networking research. In the following sections, we examine the implications of these models on future research. We group research directions into three themes: helping the network to help itself, helping users to help themselves, and helping others to help users.

Helping the Network Help Itself

There are a range of technologically-oriented directions in both HCI and networking that may facilitate more human-centered home networking.

Statistical & Self-healing Approaches

Networks that can detect and automatically fix problems offer an obvious solution to mitigating some of the difficulties of home networking. Since networks currently put intelligence in end-nodes rather than in the network itself, self-healing approaches would seem to require significant infrastructure redesign, suggesting that the Fresh Start or Gateway models would provide the most "head room" for such approaches.

Other technologies, while perhaps not providing as much benefit, may not require new infrastructure. Such technologies would exist alongside current infrastructure, much like other Bandage model technologies such as DHCP and various discovery protocols. One example of a research program in this area is the PeerPressure system [37], which uses statistical analysis to compare machine configurations and guess which one is broken. Even this system, however, requires user intervention: "only the user can recognize the sickness and therefore has to be in the loop for these steps."

Self-healing networks are not a panacea, though; they cannot understand the context and desires of users. Some aspects of network installation and maintenance may still require a human touch [8]. Additionally, when using statistical approaches to look for broken machines, privacy concerns may also be raised. For example, where is this configuration information coming from? Where is information about my network configuration being sent? There is also a danger of false positives, where highly unique configurations are flagged as broken [37].

Network Protocol Redesign

Even though there is much infrastructural inertia in the current Internet design, this does not mean that we should not explore radical options for rethinking networking; in fact, this approach is at the heart of the Fresh Start and Gateway models. Protocol redesign is a goal of a number of ongoing projects [26], and the HCI community must be involved in these. Specifically, the HCI community must provide input on how protocol design decisions can affect usability, as well as potential social impacts of the technologies, such as protection (or lack thereof) of privacy. Offering valid input to such considerations is challenging, however, as HCI as a field does not possess methods for assessing the impact of infrastructure technologies on the user experience, nor for feeding interaction design guidelines into the design of infrastructure-layer technologies [10].

Network Data Collection and Analysis

The lack of adequate network diagnostics point to the need for improved network data collection and analysis tools situated in the network itself. Such tools could record and aggregate data from various devices on the network. This information could serve a number of uses, including automatically repairing the network, (that is, helping the network to help itself), providing input for repair by end-users or other external troubleshooters, and providing the foundation for a range of tools under the Bandage or Outsource models.

Helping Users to Help Themselves

Currently, the bar to "do-it-yourself" networking is high; only a minority of householders are both motivated and skilled enough overcome this hurdle. Even though many home networking products claim to be "do-it-yourself," there are few tools available for people to manage complexity in their networks or to identify and resolve problems [34].

Visual Approaches to Networking

Since the structure and function of home networks is largely invisible, the Bandage model may suggest research efforts focused on tools and techniques for allowing end users to visually explore and manipulate their networks. Although there are many network visualization and traffic analysis tools available, they are primarily intended for use by network experts in enterprise environments rather than by network novices in the home. The majority of network traffic analysis tools give *too much* detail to be appropriate for most householders. Without a nuanced understanding of networking data structures and terminology, command-line network utilities (such as ping, traceroute, and netstat) and traffic analysis software (such as Ethereal) are incomprehensible to "mere mortals." Hence, there are opportunities for HCI community to research and create visualization systems that can provide home users with network monitoring and troubleshooting capabilities, but better align with householder needs. These tools should be designed to fit home users' conceptions and language used to describe networking, as well as their practices surrounding network troubleshooting.

Other Approaches to Configuration and Management

A range of technologies—some discussed earlier—attempt to provide new interaction techniques for device configuration and management. These systems build on existing protocols to support easier wireless provisioning [3], for example, or easier composition of services in the home [35]. Many of these tools, however, only deal with a few aspects of network configuration and management. More fully exploring the potential to layer new interaction techniques on top of or in addition to existing configuration and management protocols may be an important research direction under the Bandage and Gateway models. The Ubiquity model also suggests alternative approaches to configuration and management, including allowing householders to easily secure the individual components of their networks from unwanted access, create necessary trust associations between home devices, and ascertain the trustworthiness of other entities on the network.

Understanding Householder Motivations and Practices

Despite the prevalence of networked technology in homes, very few studies have explored why and for what purposes people want networking. Studies such as [13] represent a first step toward a deeper understanding of end user practices and social factors that make home a special place for networking – one that is unlike managed, corporate environments. No matter which model is used, understanding such motivations will give the HCI and networking communities important tools for designing new networking technology that fits the needs and practices of home users.

Helping Others to Help Users

Home users may enlist knowledgeable family and friends to help them with configuration and troubleshooting; likewise, much problem solving now is done in coordination with remote help lines. Further, the *use* of the network is collaborative, as it is an infrastructure that is both shared and adapted by all of the members of the household (and in the Ubiquity model, by all members of a community). There may be a range of fruitful research directions that explore and support the inherently collaborative aspects of networking.

Enlisting Family and Friends

Evidence suggests that householders—especially the technologically inexperienced—rely on informal networks of family and friends to solve their network problems [19, 22]. The HCI research community, however, has little understanding of householder practices surrounding technical support of home networking. The Fresh Start, Bandage, Gateway, and Ubiquity models all suggest research opportunities in understanding these practices, including: How do people seek information about how to fix their home networks? What sources do they deem trustworthy? What problems do these informal (and often reluctant) troubleshooters face? (For example, a tech savvy friend may be trustworthy, but his or her knowledge about networking may not match the user's needs [34]). The Outsource model may also suggest research opportunities focused on the social costs of calling a professional. Kiesler et al. mention that embarrassment about having broken digital devices often leads people to *not* seek technical support [19]. Are people more comfortable getting advice about how to fix a problem from their personal computer, a search engine, or a trusted friend?

Tools for Coordination

Coordinating to solve networking problems is a formidable task for computer experts [12]; hence it may be even more difficult to coordinate to solve networking problems when one does not have the vocabulary or background knowledge to explain what is wrong (which may often be the case with householders). Both the Bandage and Outsource models suggest research focused on ways to aid home users in more thoroughly sharing data with external troubleshooters. These efforts may include the development of tools that can aggregate information about the state of the network (while preserving privacy) or that train users to have a shared language for talking with experts. There is a built-in incentive for experts to install these tools and for households to use them—both want to spend less time resolving problems.

CONCLUSION

We have discussed a range of the usability issues with networking in the home, as well as the sources of many of these issues. We have described a number of possible of futures for home networking, each of which indicates different research agendas, as well as a unique set of trade-offs with respect to economics, privacy, and end-user empowerment. These futures will not be determined solely by research efforts in HCI or computer networking; rather, they will reflect the coming together of economic and social factors as well as technological factors. Despite these external factors, HCI can play a major role in shaping these outcomes in a user-centered, reflective way. As a research community, we must ask ourselves whether we truly desire a future in which users must resort either to shouldering the burden of network management themselves or enlisting paid professionals. In this version of the future, networking is merely a technological toy for those with discretionary income and significant amount of technical knowledge. Can we do better? We believe that in order to improve the usability and accessibility of networking for *all* members of society – not just the economically and educationally privileged – the HCI and networking communities need to reflect deeply on these possible models and to engage in research that can lead to a future where networking is accessible to all.

ACKNOWLEDGMENTS

We thank members of the Pixi Lab at Georgia Tech, Marshini Chetty, and our anonymous reviewers for their insightful feedback. This work is supported by grants from Linksys and NSF-CNS #0626281. One of the authors was on appointment as a US Department of Homeland Security (DHS) Fellow during this project. All opinions expressed in this paper belong to the authors and do not necessarily reflect the policies and views of these funding agencies.

REFERENCES

1. Anderson, R. Why Information Security Is Hard - An Economic Perspective. *Proc. of ACSAC* 2001, 358-365.
2. Anderson, T., Peterson, L., Shenker, S., and Turner, J. Report of NSF workshop on Overcoming Barriers to Disruptive Innovation in Networking, *National Science Foundation*, Jan 2005.
3. Balfanz, D., Durfee, G., Grinter, R.E., Smetters, D.K., and Stewart, P. Network-in-a-Box: How to Set Up a Secure Wireless Network in Under a Minute. *Proc. 13th Usenix Security Symp. 2004*, 207–221.

4. Blumenthal, M.S. and Clark, D.D. Rethinking the Design of the Internet: The End-to-End Arguments vs. The Brave New World. *ACM Transactions on Internet Technology*, 1:1, 2001, 70-109.
5. Bowker, G. Information mythology: The world of/as information. *Information Acumen: The Understanding and Use of Knowledge in Modern Business.* Routledge, 1994.
6. Brush, A.J. IT@Home: Often Best Left to Professionals. *Position paper for the CHI 2006 Workshop on IT@Home*, April 23, 2006, Montreal, Canada.
7. Cowan, R.S. *More Work for Mother: The Ironies of Household Technology from the Open Hearth to the Microwave.* Basic Books, 1983.
8. Edwards, W.K. and Grinter, R.E. At Home With Ubiquitous Computing: Seven Challenges. In *Proc.of. Ubicomp 2001*, 256-272.
9. Edwards, W.K. Discovery Systems in Ubiquitous Computing. *IEEE Pervasive Computing*, 5:2, April-June, 2006, 70-77.
10. Edwards, W.K, Bellotti, V., Dey, A.K, and Newman, M.W. Stuck in the Middle: The Challenges of User-Centered Design and Evaluation of Infrastructure. *Proc of CHI 2003.*
11. Geeksquad. http://www.geeksquad.com/
12. Goodall, J.R., Lutters, W.G., and Komlodi, A. I Know My Network: Collaboration and Expertise in Intrusion Detection. In *Proc. of CSCW 2004*, 342-345.
13. Grinter,R.E., Edwards, W.K., Newman, M.W. and Ducheneaut, N. The Work to Make a Home Network Work. In *Proc. of ECSCW*, 2005.
14. Hafner, K. and Lyon, M. *Where Wizards Stay Up Late: The Origins of the Internet.* Touchstone, 1996.
15. Hush-A-Phone Corporation and Harry C. Tuttle vs. United States of America and Federal Communications Commission, Respondents, American Telephone and Telegraph Company et al., and United States Independent Telephone Association, Intervenors. *United States Court of Appeals, F.2d 266*, 1956.
16. Hypertext Transfer Protocol – HTTP/1.1. RFC 2616. http://www.w3.org/Protocols/ rfc2616/rfc2616.html
17. Internet Protocol –Version 4. RFC 791. http://www.ietf.org/rfc/rfc0791.txt
18. Internet Protocol –Version 6. RFC 2460. http://www.ietf.org/rfc/rfc2460.txt
19. Kiesler, S., Zdaniuk,B., Lundmark, V., and Kraut, R.. Troubles with the Internet: The Dynamics of Help at Home. *Human-Computer Interaction*, 15:4, 2000, 323-351.
20. Laszlo, J., et al, V. *Home Networking: Seizing Near-Term Opportunities to Extend Connectivity to Every Room.* Jupiter Research. July 9, 2002.
21. Linksys home page. http://www.linksys.com
22. Lowe, D. *Networking for Dummies*. Wiley, 2005.
23. MacMillan, R. Plugged in: Wireless Networking Baffles Some Customers. *Reuters news report*, March 10, 2006
24. Microsoft. Windows Connect Now Technology. http://www.microsoft.com/windowsxp/using/ networking/getstarted/windowsconnectnow.mspx
25. Network Magic. http://www.networkmagic.com
26. NSF GENI Initiative. http://www.nsf.gov/cise/geni/
27. Nye, D.E. *Technology Matters: Questions to Live With.* MIT Press, 2006.
28. Rose, B. Home Networks: A Standards Perspective. *IEEE Communications Magazine*, Dec 2001, 78-85.
29. Rosenberg, N. *Exploring the Black Box: Technology, Economics, and History*. Cambridge U. Press, 1994.
30. Scherf, K. Consumer Electronics Association (CEA) Conference, Parks Associates Panel on Home Networking, 2002.
31. Shehan, E. and Edwards, W.K. Pinning the Tail on the Networked Donkey: Why IT@Home Needs Network Visualization. *Position paper for the CHI 2006 Workshop on IT@Home*, April 23, 2006, Montreal, Canada.
32. Standage, T. *The Victorian Internet.* Berkley, 1998.
33. Star, S.L. The Ethnography of Infrastructure. *American Behavioral Scientist*. 43: 3, 1999
34. Teger, S. and Waks, D. End-User Perspectives on Home Networking. *IEEE Communications Magazine*, April 2002, 114-119.
35. Truong, K., Huang, E., and Abowd, G. CAMP: A Magnetic Poetry Interface for End-User Programming of Capture Applications for the Home. *Proc. Of Ubicomp 2004*, 143-160.
36. Wang, Harry. “Networks in the Home: Connected Consumer Electronics.” Parks Associates, 2006.
37. Wang, H.J, Platt, J.C., Chen, Y., Zhang, R. and Wang, Y.-M. Automatic Misconfiguration Troubleshooting with PeerPressure. *Proc. USENIX OSDI*, 2004.
38. Wang, Y.-M., Russell, W., Arora, A., Xu, J. and Jagannathan, R. Towards dependable home networking: An experience report. *International Conference on Dependable Systems and Networks*, 2000.
39. WiMAX Forum. Mobile WiMAX: The Best Personal Broadband Experience! 2006. http://www.intel.com/n etcomms/ technologies/wimax/Mobile_WiMAX_WP.pdf