

DOI:10.1145/1953122.1953143

Before building the network or its components, first understand the home and the behavior of its human inhabitants.

BY W. KEITH EDWARDS, REBECCA E. GRINTER, RATUL MAHAJAN, AND DAVID WETHERALL

Advancing the State of Home Networking

THE CREATORS OF the original Internet architecture imagined a network of networks linked by protocols connecting the world's diverse and disparate networks. That vision has been profoundly successful, with the Internet reaching beyond the government and corporate contexts in which it was born into settings ranging from public hotspots to places of worship, to rural sub-Saharan Africa, even to space.

In the Western industrialized world, networking is rapidly being adopted in the home; for example, in the U.S. in 2009, approximately 63% of homes had a broadband connection, and over 50% had a "home network," defined as multiple computers sharing a broadband connection via either a wired or wireless network within the home.¹⁷ These networks are part of the global Internet, participating in it just like corporate, government, and other networks and representing an increasingly large percentage of nodes on the global Internet.

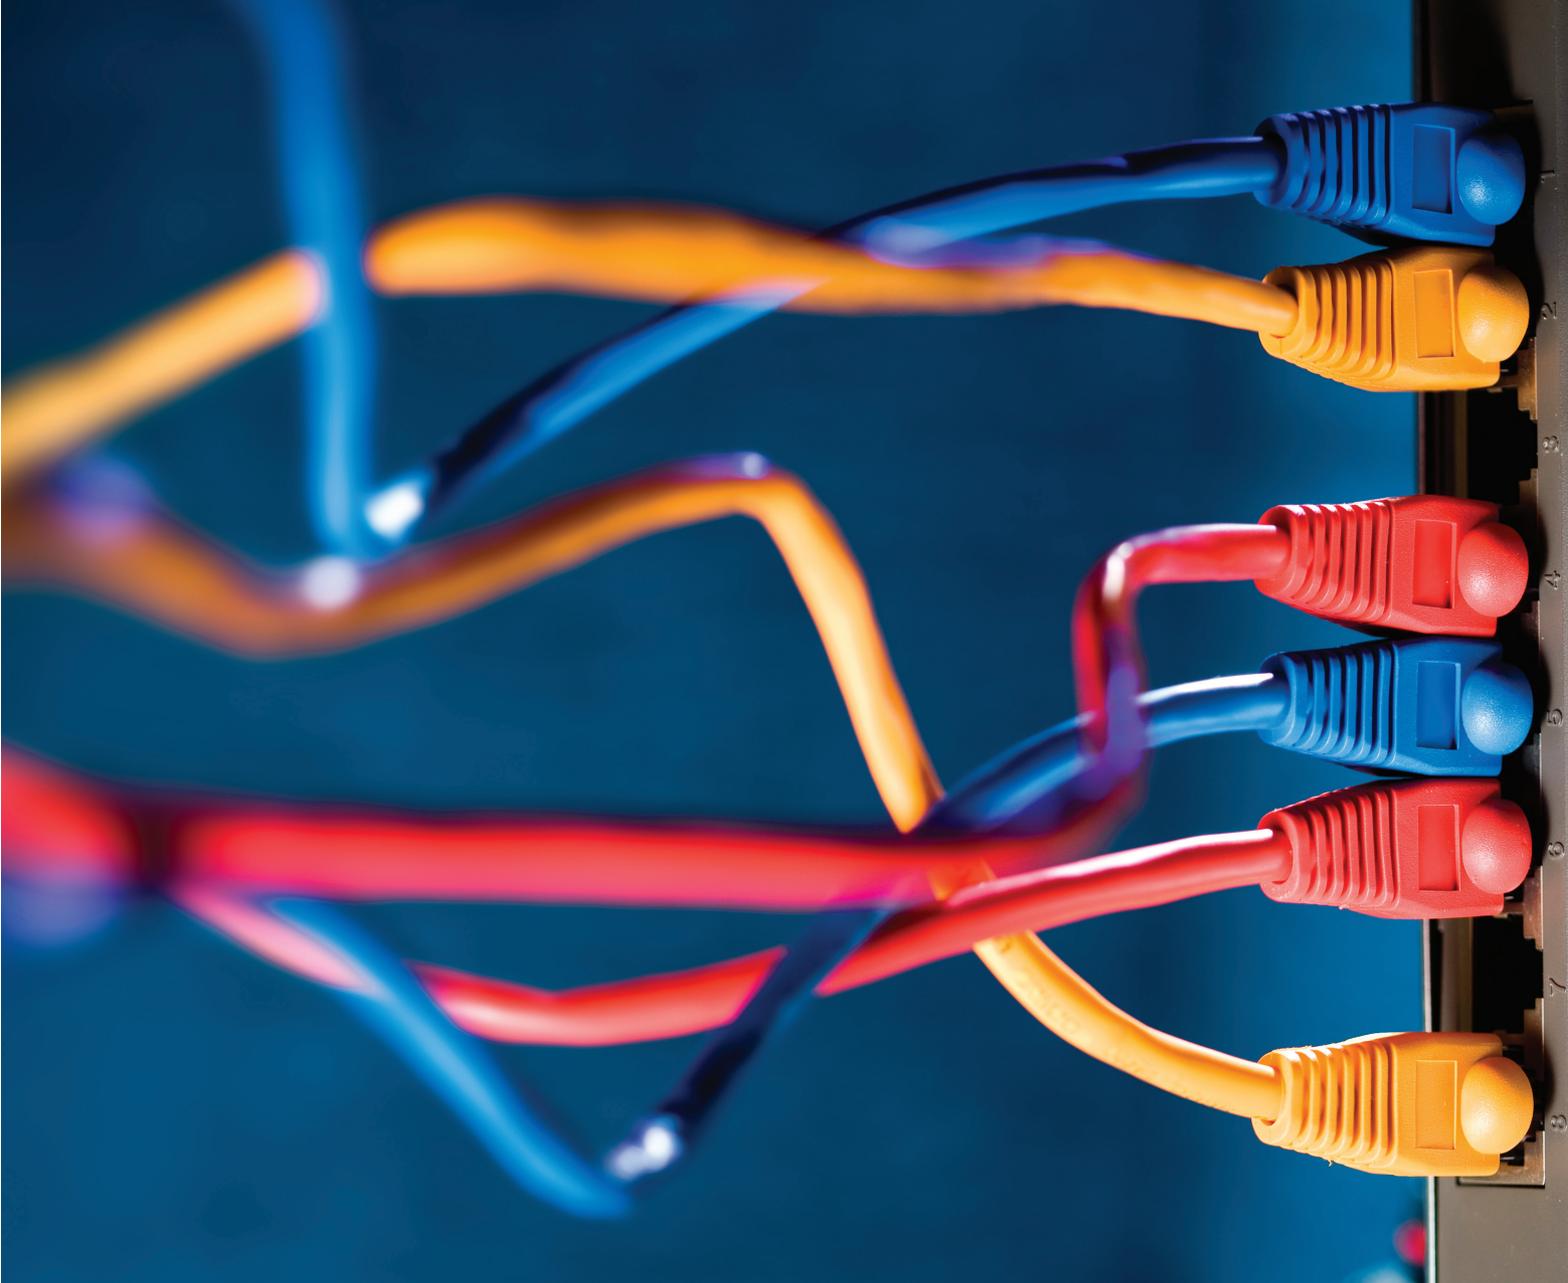
By such measures, home networking is a success. The potential benefit of this deployment is huge because it opens the home to new commercial services in entertainment, education, health care, and communication. These applications have positive implications for sustainability (such as through telework), public access to governmental services, and better care of elders in their own homes. All these applications depend on secure, manageable, cost-effective deployment in the home.

However, home networks are fraught with problems, particularly in terms of the difficulties end users face managing and securing their networks.^{8,9} Many of these problems arise because Internet-style networks, which were developed for managed environments, have been transplanted relatively unchanged to the home. In the best cases, the problems result in mere hassle and frustration for householders; in the worst cases, they pose a threat to individual privacy and lost opportunity to adopt next-generation applications. For example, while advanced home-automation technologies (such as those by Crestron and Control4) have been available for more than 20 years, adoption is limited, with fewer than 159,000 units shipping worldwide in 2009.¹

As a starting point for fixing problems with home networks, we articulate

» key insights

- **Usability of home-networking technologies is a key impediment to adoption of new applications in the home.**
- **Network usability problems run deep because the technology was originally developed for research labs and enterprise networks and does not account for the unique characteristics of the home: lack of professional administrators, deep heterogeneity, and expectations of privacy.**
- **Addressing the challenges of home networking is not simply a matter of designing better user interfaces but a concerted, interdisciplinary effort by networking, HCI, and systems researchers.**



the diverse underlying factors responsible for these problems, rooted in technical aspects of the Internet architecture and protocols, in human-oriented aspects of householders and the home setting itself, and in the economic and commercial factors hindering effective solutions to the problem, even though the market is lucrative.

This diversity of factors makes clear that the problems in home networking today are not due to core technical issues nor exclusively to issues that can be solved through better user education or user interfaces (UIs). Dramatic improvement is unlikely to be brought about by addressing the two sets of issues independently, whether by painting a UI veneer on top of existing technologies or by designing new protocols and architectures without accommodating human practices,

needs, and routines. There is a deep connection between the user experience and the network's underlying capabilities, leading us to suggest the two sides of the problem must be tackled together through cross-disciplinary research approaches.

Additionally, such approaches must be developed bearing in mind the uniqueness of the home environment—lack of trained administrators, extreme heterogeneity across homes, and strong privacy considerations. A large body of work addresses how to simplify network management and diagnosis in enterprise networks and Internet service providers (ISPs), but the considerations at play in the home mean it cannot be applied directly or easily adapted there.

We conclude by sketching example approaches combining technical capa-

bilities and sensitivity to the user experience while being compatible with the home environment. We do not claim these approaches are the only fruitful directions but hope they will stimulate the research community to try fresh cross-disciplinary approaches.

Strains on the Home Network

To understand the home network, we should appreciate where it comes from:

The Internet went home. Internet use in the home has followed a path similar to many other advanced home technologies. Tech-savvy “early adopters” were the first to bring networks into their homes, often to support telework.²⁶ They had both the technical sophistication and the hobbyist's motivation to persevere. In the 1980s and earlier, installing a home network

meant being committed to installing and managing commercial-grade routers and switches, running structured wiring throughout the home, and explaining to the local telephone carrier that one needed, say, a T-1 or ISDN connection.

Internet adoption in the home went mainstream following the advent of the Web, with the standard Internet Protocol suite deployed in consumer operating systems like Windows 95. The growing number of Web sites attracted consumers who had previously participated only in closed, paid online services like AOL and CompuServe. This period brought about the consumer ISP, as well as a commercial market for consumer-grade networking hardware. The role of home networks during this period was generally limited to sharing the Internet connection with multiple computers (such as one in the home office and another for children) and potentially providing access to a shared printer.

The next big shift in home networks occurred around 2001, as inexpensive wireless connectivity in the form of IEEE 802.11b made it possible to provide high-speed network connectivity throughout a home without structured wiring. Home networks today have shifted beyond simple Internet and printer sharing to include new devices and applications (such as media streaming, game consoles, and

WiFi-enabled phones), as well as new connectivity options (such as faster versions of WiFi and Ethernet over powerlines), helping spread the network throughout the home.

Figure 1 is from a Web site dedicated to helping hobbyist users show off their network topologies, depicting an advanced home network. Note it includes a mix of wireless and wired connectivity, along with a range of devices, including PCs, printers, smartphones, cameras, and game consoles.

The broad strokes of Internet adoption in the home are well known; we lay them out to highlight that when network adoption in the home shifted from a trickle to a flood, the network brought home was the same Internet adopted in the corporate world and designed for environments radically different from the home.

The home is different. The benefits of adopting the “one true Internet” at home are profound, as well as obvious, benefitting consumers through standardization on widely adopted, open protocols; applications and devices built for the Internet protocol suite can be run at home as easily as they can in the enterprise; and economies of scale allow ever more inexpensive networked devices to find their way into homes, as in the figure.

At the same time, adoption of the Internet in the home brings problems due to the misfit of the technology with

the home’s dynamics and context. The Internet architecture was developed for a world of technical sophistication (where experts handle network connectivity and management) and shared trust (where network operators share responsibility for protecting the network’s integrity). These properties do not carry over to home networks.^{4,23}

Home networks differ from other Internet-style networks (such as enterprise networks, data centers, and ISPs) in three significant ways:

No professional administrator. Home networks do not have professional, trained “administrators” in the same sense as other networks,³ making difficult their management, security, and diagnosis;

Deep heterogeneity. The home network is a place of deep heterogeneity, even experimentation, as new applications and devices are installed, and householders push the technological boundaries of their networks. This heterogeneity means applications and services in the home may be deployed onto new and potentially unforeseen home-network infrastructure and configurations, and conversely that the network infrastructure in the home must be able to support new types of applications and services introduced by users. This complexity means isolating any resulting problems may be difficult; and

Expectation of privacy. There is an expectation of privacy in the home that is different from what is expected in the workplace. Whereas users in most companies have no expectation of privacy on their networks (due to policy or the needs of the IT department), home users have a strong expectation of privacy, extending even to restricting certain access to those who help manage them (such as ISPs and professional troubleshooters). Also, unlike corporate networks, homes are unlikely to have centralized, managed-access controls for their devices and information.

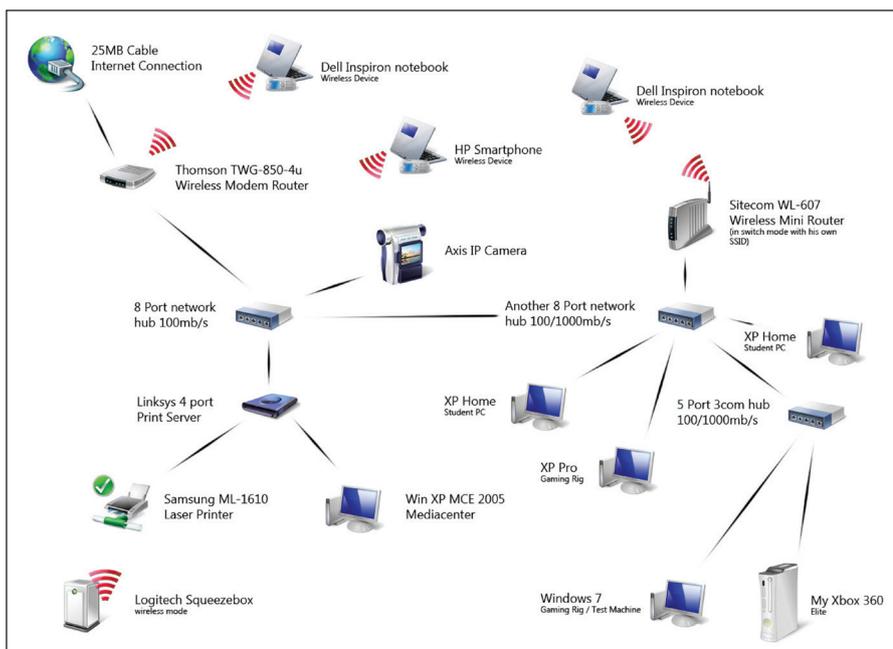


Figure 1. Schematic of real home network.

a We deliberately do not say home users “are unsophisticated.” Rather, home users are deeply sophisticated about their individual domestic practices and the routines of their own homes, along with how home networking should be meshed with these practices and routines, about which the purveyors of home networking solutions are conversely generally unsophisticated.

Rising tide of problems. These factors diminish the ease of using the home network for householders. Heterogeneity may increase the number of network faults and difficulty of isolating problems. Users often have neither the technical training nor the motivation to deal with them, raising frustration and support costs for service providers and vendors. Beyond these factors, the need for privacy complicates solutions that may be workable in the enterprise (such as remote administration by paid professionals).

It is an exaggeration to say that home networks are broken, since they have proliferated, and many people are able to watch Internet movies at home and back up their photos. However, there is also a significant tax on user goodwill due to the time and effort spent (wasted) dealing with home-networking issues; for example, in 2008, approximately half of all home technology users in the U.S. needed help from others to set up or install new devices.¹⁸

The state of home networking clearly results in user frustration; it also translates into significant monetary costs for companies. For example, in 2002 when home networking was relatively new, Parks and Associates²² found that home network devices were the single most returned category of items at “big box” consumer electronics stores in the U.S. As recently as 2010, over 25% of all wireless access points were returned, despite most presumably functioning as intended.¹⁹ Data indicates only 5% of consumer-electronics product returns worldwide are due to actual technical failure.²

Furthermore, the current state of home networks functions as a barrier to adoption of new technology (such as home-automation systems). Indeed, consumers in the U.S. regularly cite technical complexity as the primary disincentive to adding new networked devices in the home.¹³

The silver lining, if there is one, is that home networks also represent an opportunity for research with benefits beyond merely improving the user experience, increasing security, or decreasing support costs. If computer scientists can advance the state of home networking, then they may enable new applications to move out of the lab into

practice. Many of the applications envisioned by the research community—from entertainment and social media to telework and remote collaboration to health and wellness—require functioning, manageable, correctly configured, secure home networks. Removing the barriers to such networks has the potential to unleash a wave of innovation in the networked home.

Fault Lines

Given the problems around home networks, the natural question for computer scientists to ask is whether they are easily fixed; for example, does home networking technology simply suffer from a lack of UI polish? Indeed, this is not the case, and the underlying causes are deeper and structural. Some causes stem from technical aspects of the Internet architecture and protocols, others from human-oriented aspects of householders and the home, and still others from economic and commercial aspects. Together, the picture they paint is problematic and unlikely to improve greatly without sustained involvement from multiple computing disciplines.

Technical. The technical causes are, by and large, an inherent aspect of the design of the Internet architecture:

Necessity of configuration. A core precept of the end-to-end argument espoused as a key architectural principle of the Internet is that the network itself is an application-neutral carrier of bits; on the other hand, end nodes must be able to speak a wide (and potentially open-ended) array of application-layer protocols. End-user devices are often programmable and general-purpose and must be configured correctly to work on the network. In the home, end users generally perform this configuration, possibly involving settings at the link layer (such as network name and encryption keys), network layer (such as default gateway, Domain Name Service, or DNS, settings, and IP addresses), and application layer (such as Web proxies and print servers). While some of these settings may be handled by network services (such as Dynamic Host Configuration Protocol, or DHCP), many more require direct user intervention to be set correctly.

Note that the need for configuration is not an inherent characteristic

of communication networks. For example, in the public switched telephone network (PSTN), a different set of design decisions manifest in a radically different user experience, where infrastructure devices and end-user device settings are entirely removed from the purview of end users. A landline phone plugged into an RJ-11 connector “knows” its phone number and is immediately integrated into the global network without the user having to be configured in any way. On the flip side, such a device does not support the radical extensibility to new applications or even the entirely new protocols many Internet devices have today.

But where configuration is a necessity there is also the possibility for misconfiguration, possibly exposing the network to privacy and security risks or even to the catastrophic severing of a device’s connectivity. Exposing settings to users and generally requiring that they interact with them presents a distinctly unfriendly user experience in which users must grapple with technical jargon and low-level settings far removed from what they want to do. This configuration task also tends to be highly complex due to several factors:

Interaction of applications and network. Application developers tend to think in terms of an ideal world where devices or applications abstractly “sit on top of the network,” viewing it as a more-or-less opaque infrastructure that carries bits. Unfortunately, this abstraction does not match reality, and the mismatch is sometimes exposed in application failures difficult for users to understand and fix.

Consider that homes today provide great topological flexibility, possibly including multiple wireless access points, powerline Ethernet bridges based on the HomePlug standard, structured wiring, and Ethernet-over-coaxial solutions (such as the Multimedia Over Coax, or MOCA, standard), each with its own adapters and configuration utilities. This infrastructure does not always work with applications unless configured in highly specific, technical ways; for example, adding a new access point may introduce a second DHCP server and subnet, resulting in some clients getting IP addresses incompatible with the rest of the home

network. Multicast discovery protocols may then be unable to cross subnet boundaries. The result is that some applications may stop working properly when new and seemingly unrelated devices are added.

Applications may also require the infrastructure be configured in specific ways; for example, playing an Internet-connected game may require tweaking firewall or Network Address Translation (NAT) settings so the game communicates with outside servers. While relatively new protocols (such as the NAT Port Mapping Protocol, or NAT-PMP) alleviate some of this manual tweaking by enabling programmatic configuration of NAT port forwarding, many manual tasks persist; for example, sharing a folder of pictures with a distant relative may require users set up Dynamic DNS (so the relatives do not need to know the potentially ever-changing IP address of the home router), understand static-versus-dynamic addressing on the home network, or change network and host security settings. Such tasks require users to step outside the application to modify the network; these tasks are unlikely to be easily automated in the near future because they crosscut multiple layers of the networking stack, as well as multiple applications and services.

Today's home networks do not provide good abstractions to applications. End users must manually configure the network to implement their goals. To do so, they are essentially required to be network managers for a small subnet within the Internet, managing their devices, network infrastructure, and applications.

Conflation of policy with mechanism. Further complicating the task of managing home networks is the fact that the controls available today are generally low-level and divorced from any high-level policies users may wish to establish for their networks. These controls (such as to tweak routing and translating and forwarding packets on the home network) and parameters (such as the size of the maximum transmission unit, or MTU, and wireless security primitives) are difficult to translate into actionable, high-level policies at the user level.

This tension is perhaps most apparent in network security, in which

security policies are deeply intertwined with low-level network mechanisms and topology. The situation today is characterized by devices on the “inside” of the network having an implicit trust relationship with one another, along with often weak host defenses. This means network misconfiguration can potentially expose poorly defended hosts to a range of attacks. Further, the conflation of security with relatively coarse-grain configuration controls makes it difficult or impossible to express certain security goals using the set of available “knobs.”

As an example, consider how one might securely allow a visitor to access one's home network. This secure access is handled in most home networks in one of two ways: turn off wireless security completely (exposing all hosts on the network to public access) or give the wireless key to the visitor. Even giving the key to a visitor—undeniably better than turning off security entirely—may not map well to the user's intended policies. In most home networks, having access to the network itself includes not only the shared Internet connection but all the hosts on the network. While this may be acceptable for some family members and close friends it may be completely unacceptable for visiting technicians and friends of teenage children. Some newer wireless routers offer a separate guest wireless network, allowing guests to access the Internet but not the other devices connected to the home network. While this feature represents a step in the right direction, it is too coarse-grain; for example, it does not support users wanting to let their guests use the printers in the home but not access media files.

Human-centric. Not all problematic aspects of home networking are directly or solely technical in nature; many interact with human needs, behavior, and expectations. Here, we outline some of the causes of home-network problems rooted in these human concerns.

Paucity of conceptual models. Home networks do not provide solid “hooks” for users to form reliable, predictable conceptual models of how the network functions. The importance of such models has been long known (Norman¹⁵ is a classic example), allowing

users to observe the behavior of a potentially complex system and formulate a plausible model for how it works based on their observation, using it to control the system. Such models need not correspond completely with “reality” to be useful. Kempton's work on “folk theories” of home-heating controls¹² illustrates how even incorrect models allow people to work with complex systems.

However, current networks do not expose the kind of information necessary for users to form such models; for instance, much of the behavior and state of the network is unobservable, as are many effects of changes on the network. For most users, the only readily observable states are whether the network is working and whether it is dysfunctional. Wireless technologies exacerbate the problem, as there is not even a cable to show how data might flow. The effects of a lack of conceptual models show up almost immediately in user-oriented studies of networking,^{8,9} whereby ordinary consumers are often baffled by network behavior, unable to devise a troubleshooting strategy that is more sophisticated than hopeful rebooting and often unaware of even what capabilities are provided by their own network. Figure 2 (from Grinter et al.⁹) is in stark contrast to Figure 1, showing an example of a household's sketch of a home network and illustrating the lack of an actionable conceptual model; note the absence of any aspect of network topology, link type, or router.

Broken expectations. Desire for communication, entertainment, and other applications is the key reason households adopt home networking in the first place. Indeed, desire for relatively advanced functionality (such as in-home media sharing among home devices and media sharing across homes) factor prominently in home-network adoption.⁶ However, achieving such functionality in a way that fits the ecosystem of the home is beyond the grasp of many, betraying the expectations and desires that lead users to adopt home networking.

Part of the problem is that households' nontechnical requirements are often not addressed by home-network technologies. Requirements like device aesthetics, flexibility of device

placement, and a device's ability to be used by multiple members of the household are often highly relevant in the domestic setting but rarely determined solely by a device's technical properties. Even though there may be nothing technically wrong with a device, it may still not fit well in the domestic ecology of the home due to social, cultural, aesthetic, or other constraints.

A second aspect of the problem is the expertise required to install and use advanced home-network devices. Device- and network-management chores are often divided among householders, meaning not only do new devices bring new domestic labor into the home but the expertise needed to fix problems may be spread among multiple members of the household or even among family and friends outside the household. Thus, the technical burden of keeping the network running becomes a social burden, whereby householders must leverage family and friends for support.

Personalized home network. Any effort to make networks more usable is immediately complicated by the fact that the home network, as with any domestic technology, is deeply intertwined in the everyday social routines and realities of home life. Much prior research has demonstrated the degree to which home networks reflect household domestic routines; see, for example, Tolmie et al.²⁵ The demands of a particular household may govern how and when the network is used, the placement of machines (to allow parents to keep an eye on children's Internet use or for aesthetic reasons), what applications are present, and what media or content are considered "shared" by the family versus owned by a particular individual. These human considerations, in turn, drive technical heterogeneity in the network; for example, a desired device placement could motivate installation of such new infrastructure devices as wireless and HomePlug; such desired applications as media sharing could lead to introduction of new devices onto the home network; and the patterns of sharing and privacy in a given home may dictate aspects of network configuration, device configuration, and content placement. Moreover, there may be as

many variations along each of these dimensions as there are homes.

The deep personalization in today's networks is unlikely to go away over time; studies of other domestic technologies illustrate the degree to which the domestic order varies across households.¹⁶ Such personalization also complicates naïve approaches to "solving" the problems of home networking; many of us have called tech support only to find the "expert" at the other end of the line has no knowledge of the particulars of the home network and is simply following a scripted troubleshooting routine. One core issue with such approaches is that the remote troubleshooter is unaware of any local customizations or idiosyncratic uses of the home network; such a one-size-fits-all approach is likely to remain problematic for the home, given its deep personalization.²¹

Privacy and the home network. Home networks are not only deeply personalized but also deeply personal. Users have information on the network they do not want to share with others outside the home and in some cases even with those within the same home. The expectation of privacy within the home runs deep, and is different from

any such expectations that may exist in most corporate networks. These differences can also complicate the task of seeking help when the network does not function as desired. Many home users might, for instance, be uncomfortable giving unfettered access to their entire network and the information on it to an outside technician, even if the technician is there to help repair the network.

Economic. One might argue that, given the size of the market for networked devices and services in the home, market pressure would drive commercial interest to fix the problems of home networking. Commercial players are certainly tackling problems; for example, Cisco's Valet is a recent attempt to simplify the setup for wireless routers,¹⁹ and Apple's Bonjour enables easier discovery of devices on home networks.²⁴ Such solutions, while beneficial, are still too limited to truly advance the state of the art. Several economic factors tend to damp more sweeping but potentially more effective long-term change, either unilaterally on the part of companies or as part of a broader effort.

Lack of data. Solving a problem requires reliable data, but data on prob-

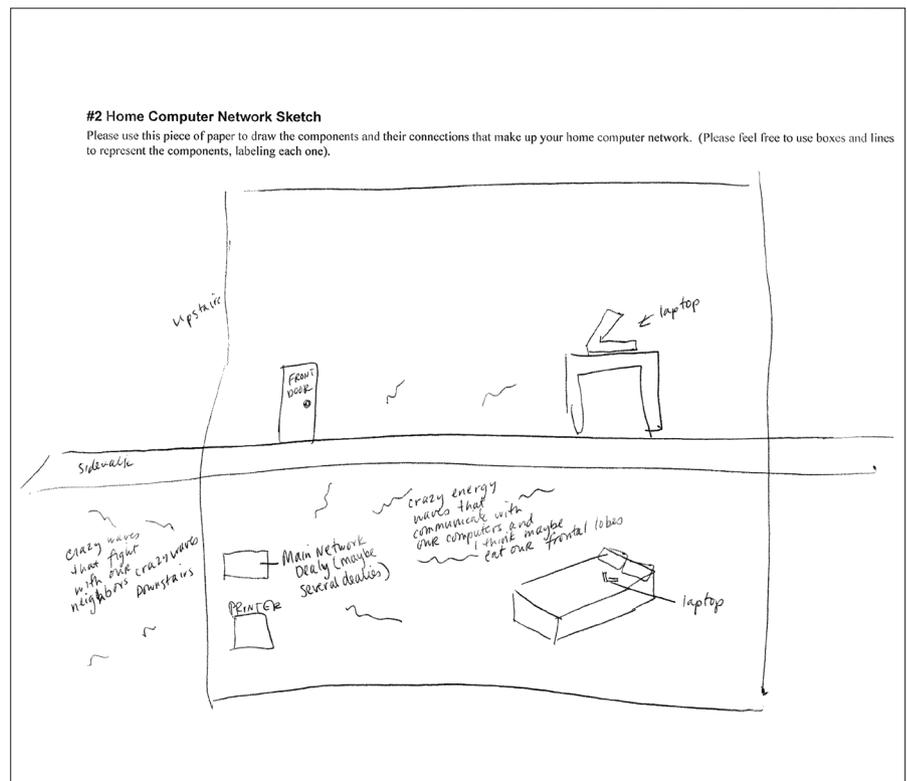


Figure 2. Sketch of a home network by a householder, who, though a regular user of the network (set up by someone else), was incapable of troubleshooting or extending it.

lems inside home networks is difficult to come by. Since each home is different, one needs data from numerous homes to address the problem broadly. Ethnographic studies (such as those cited earlier) provide rich accounts of how individuals understand and work with a network within the social setting of the home. Yet they cannot provide detailed technical data about what is actually happening on home networks, nor can they be widely scaled to understand how network-level behaviors may vary across households.

Automated problem reporting would be valuable, in much the same way the introduction of Microsoft Windows's error reporting led to significant improvement in bug fixing.⁷ However, databases of user reports of home-network problems are much less useful today, spread across multiple organizations, including ISPs, operating system vendors, and device vendors. This is because users tend to report problems to different organizations, depending on which they perceive (correctly or incorrectly) as responsible. Thus, each organization has a limited window into the problems while also being reluctant to share for competitive reasons, and no one party has a holistic view that makes it easier to develop effective manageability solutions for the home. Further, user reports of network problems may not correspond well to the actual state of the network; a call to a technical support line saying "The network doesn't work," followed by instructions to reboot the router and all attached devices, does not reveal the actual source of the error. While self-reports, call logs, and trouble tickets may describe a user's perception of a problem, they do not represent ground truth about the network itself.

Cost of heterogeneity. A complication that makes it costly to develop services for the home is the high degree of diversity across homes, with deep personalization leading to this heterogeneity. Homes today have different sets of devices, topologies, and user preferences that will likely always be the case. Diversity leads to many possible network configurations, as well as failure symptoms and root causes. For ease of use, network products and services must tackle this diversity, despite increasing the complexity of use, along

with the costs of development, testing, and support.

Evidence of the difficulty of handling diversity across home networks is the increase in vertical integration of hardware and software; for instance, security-products manufacturer Schlage (<http://www.schlage.com/>) offers home-automation devices (such as door locks and lamp modules). Not only do they not integrate well with other vendors' devices, to enable remote access to them, users must acquire an additional Schlage device (called LiNk). While this architecture makes the solution easier to develop and more reliable by constraining the problem, it ultimately increases costs for providers and users alike, and the lack of device composition limits functionality.

Standard interfaces for cross-device interactions may eventually mitigate the problem. However, progress has been slow at best due to the lack of clarity on the specific problems and the tussle in the marketplace as companies seek to differentiate their products. It is not in the business interest of a particular vendor to make a significant engineering investment that might lift the market as a whole.

Burden of support. Even if home networks become much easier to use, some faults will still occur. When something goes wrong that users cannot remedy easily, their recourse is to call the service provider for technical support or return the product. Diagnosing and repairing these faults is expensive for vendors in today's cost structure. Monthly service charges are small relative to the cost of qualified technical personnel; even a single service call can wipe out the profit due to a customer. Returns often represent the loss of a customer and all associated revenue rather than an actual equipment fault that can be repaired. Both factors make home networking less lucrative than it might otherwise be.

A more perverse aspect of support costs is that they are not always borne by the responsible party. From the householders' point of view, the network has "gone down" if they are unable to communicate through a networked application, easily leading to the wrong party receiving a demand for technical assistance. A classic example is when computer problems are

misconstrued as network problems, as when a computer virus severs network connectivity. Users mistaking the problem are hardly to blame; some studies have shown that users deal with from three to seven service providers, including ISPs, cellular data plans, email providers, and others, on average in the U.S. to "make the network work," in addition to multiple hardware and software vendors.⁹ Faults may lie with the computer, network infrastructure in the home, ISP, or remote servers, and there are no well-established attribution mechanisms. The result is to dump costs in ways that complicate advances in home networking.

Research Agenda

These fault lines help explain why home networks are difficult to use and secure and expensive to support but can be reinterpreted as a research agenda to advance home networks and the ways they are used. Pursuing it is not simply about reducing the costs of the status quo but about enabling significant innovation. After all, if users are unwilling to install a piece of equipment for fear of breaking the network, how will computer scientists enable novel applications in the home for health care, entertainment, sustainability, or other worthwhile needs?

Transforming the fault lines into a research agenda calls for action to address each factor we've identified. To succeed, the research requires two properties: The first is it must specifically target the unique characteristics of home networks: lack of trained administrators, high level of heterogeneity across home networks, and user expectations of privacy. The home is a drastically different setting from the enterprises and government labs that developed and first adopted the Internet.

This difference requires researchers to reexamine aspects of our accepted wisdom to truly support this growing portion of the global network, and means approaches for network management and diagnosis developed for other settings (such as enterprise networks¹¹) are unlikely to apply directly to the home network. For example, some tasks (such as strict reachability limits for machines that hold payroll data) are not critical in the home, while others (such as managing consumer-

grade devices that lack support for enforcing policies) are indeed unique to the home. It is likely that some of the underlying information-processing techniques of enterprise tools (such as those for failure correlation and localization) are usable in the home. But even in these cases, the focus of the techniques is often on scale. In the home, the focus must be on making results accessible to end users. The results of today's diagnostic systems are often difficult to interpret even for trained professionals.¹⁴

The starting point for effective management research is often data that provides insight into real problems. However, the uniqueness of the home setting means existing data sets—collected in other networking contexts—may not be applicable. Even though computer scientists may have access to great volumes of data about traffic on the backbone Internet or routing behavior internal to enterprises, they have shockingly little insight into what happens in home networks. Ethnographic studies have begun to reveal the human side of the equation, but they must be coupled with data about the network itself. This goal also represents a technological challenge, and one might imagine different approaches for collecting and correlating such data; for example, the combined data could be supported by configuration-reporting mechanisms that would let users send a summary of their network state, along with a problem report. However, developing such facilities is difficult because the configuration data must be extracted from a distributed, heterogeneous set of devices at a time when the network may not be working. Whatever method for collecting the data is considered ideal, there is value in acknowledging that the uniqueness of the home makes it a necessary (and worthy) site for further data collection.

Second, research that rethinks the management of the home network must involve co-design of the networking and human-computer interface (HCI) aspects of any solution. Many of the problems in today's home networking straddle both human and technical challenges, so are neither solvable by computer systems and networking researchers focusing in



Exposing settings to users and generally requiring that they interact with them presents a distinctly unfriendly user experience far removed from what they want to do.



isolation on new technology nor by HCI researchers focusing in isolation on user-experience approaches. These two sides of the problem must be tackled hand-in-hand because there is a deep connection between the desired user experience and the network's underlying capabilities. Rethinking the home network requires more than HCI researchers painting a veneer of UI on top of existing, unworkable technologies; likewise, it requires more than networking researchers designing new protocols and architectures in the absence of knowledge about human practices, needs, and routines in the home. For example, how data is collected might first appear to be a purely technical challenge but touches on user perception of privacy. Similarly, providing more realistic conceptual models of the network (perhaps through end-user-centric visualization tools) may drive new technical requirements for specific forms of instrumentation and monitoring in the home network; developing automated approaches to troubleshooting must be done with understanding and respect for the variance across households in terms of domestic routines.

Here, we offer four examples of potential projects that cross the disciplinary divide between networking and HCI, intending not so much to chart a complete trajectory for interdisciplinary research in home networking but to illustrate the kind of interaction that may result from such collaboration:

Privacy-respecting remote diagnosis. Remote management systems (such as Intel's Remote PC Assist Technology, or RPAT) are emerging,¹⁰ with related tools and services allowing a remote operator to directly access and reconfigure a computer. This remote technology can lead to significant improvement when users act as agents of change while on a support call. However, when used in the home, the technology raises privacy concerns due to the private data on home computers and the traditional situation of an administrator having unfettered access; technology from the corporate world may not be the best fit with the home context.

How might ISPs provide remote-diagnosis tools that support the home's unique privacy constraints? Clearly,

it is not necessary for every remote-support technician to be able to view any and all contents to diagnose and repair network connectivity problems nor retain configuration data that permits future access. Creating remote-diagnosis tools for the home involves a synthesis of technical- and human-oriented work involving understanding users' orientation toward privacy (what information can be revealed and when) and accountability, allowing users to see what remote users have accessed and changed, potentially even rolling back these changes if unhelpful. It also involves creating new mechanisms to allow data to be collected from disparate elements of the distributed system that is the home network, security techniques to ensure that protected user information is not revealed, and frameworks that allow reversal of changes to networks or hosts as needed.

Leveraging social networks. Most of us are familiar with at least one effective technique for coping with the complexities of home networks: go to family and friends rather than wait on the help line for expert guidance. Householders use it not just for cost or availability reasons but because family and friends are much more likely to understand the complex, personalized, situated nature of our home networks and home routines.²⁰ Home networks could be much more effective at supporting such collaborative troubleshooting; they provide no support for it today. Imagine tools that extract a summary of the home network configuration and search a householder's social network for friends with similar setups or even with shared problems. These are the friends most likely to be able to help. Some research systems, notably the NetPrints system,³ come close to exploring this approach by collecting collaborative databases of network problems to diagnose problems and reveal possible solutions. As with remote diagnosis, privacy is a chief challenge in exposing enough of the network internals to facilitate troubleshooting without leaking information users might consider sensitive.

Leveraging the cloud. Professionally managed cloud-based services may play a key role in home networks of the future. Indeed, even today the



Application developers tend to think in terms of an ideal world where devices or applications abstractly “sit on top of the network,” viewing it as a more-or-less opaque infrastructure that carries bits.



cloud provides an opportunity to circumvent tasks that would be exceedingly complex in home networks; for example, uploading a folder of photos to a service (such as Flickr) to share with family and friends is far easier than configuring the home network to share the folder directly. In effect, the cloud provides an opportunity to outsource certain services to a professional (often for-profit) provider, removing the management burden from householders. The movement of users' personal data—email, calendars, photos, music libraries, even financial data—into the cloud is increasing. However, as with other forms of outsourcing, privacy risks might be associated with this movement.

What about future research and commercial opportunities? The cloud may enable a shift in other sorts of management tasks; for example, lightweight, potentially simpler client devices may depend more on the cloud for services (such as management and data backup). Ubiquitous, low-cost, wide-area wireless may even make it possible for the home network as a separate, discrete entity to vanish, replaced by direct connections from each device to the cloud. In each case, however, the challenges are not just how to ensure adequate connectivity and robust security. They also touch on humans being able to understand the implications of the shift and manage their relationship (and their data's relationship) to the services being offered and that they pay for.

High-level attribution tools. Standard tools that can attribute a problem with the operation of a networked device in the home to a responsible party would lead to a more direct and efficient path to problem resolution. Simply being confident that a problem is with the wireless access point rather than the ISP is a major step forward. However, to be effective, such tools should produce results that map to the conceptual models of users, especially since so much problem solving requires shared agency between tools and the people using them, with both working together to identify and correct problems. This style of interaction is very different from what exists today. Testing tools (such as ping and traceroute) pro-

vide low-level results, even if they are covered with a UI veneer as part of a troubleshooting “wizard.” These results may not be helpful for attribution; for example, if a computer fails to connect to an access point, the user may not know whether the problem is with the computer or with the access point. New tools that identify not just symptoms but also potential causes of network problems must be integrated with interfaces that allow users to assess the network, test hypotheses, and supply information that may be unavailable to the tools themselves.

Balancing self-configuration and user control. Methods for engineering networks that configure themselves would be valuable for reducing the user’s burden, as well as the possibility of misconfigurations. This might be accomplished in a number of ways, each with its own trade-offs; for example, it might be possible to impose a fixed topology (such as requiring all devices in the home connect to a single centralized gateway) to limit the range of faults that might occur and impose a single point at which policy is enforced.⁵ Other approaches might leverage the cloud as an intermediary by, say, serving as an external rendezvous point or by outsourcing services (such as storage) to a place where they are more easily configured by technical professionals.

However, a fully self-configuring network, removing all customizability from user control, is unlikely to be workable. Given the level of personalization and desire for integrating the network into domestic practices, users will always need some form of control. The challenge is finding and striking the right balance between self-configuration and user control. What things can be removed from users’ purview, and over what things must users retain control? Computer scientists do not yet know enough about the right balance to begin to inform technical solutions to self-configuration.

Conclusion

While adoption of home networks has steadily increased since the late 1990s and early 2000s, this growth also reflects deep problems and limitations. The problems are a result of the colli-

sion between technical design choices and fundamental aspects of the human condition. The problems millions of users now face with their networks go beyond inconvenience and nuisance. Threatening the privacy of home users, they also pose significant barriers to the adoption of next-generation applications in the home in areas as diverse as health care, sustainability, and education. Addressing them is inherently cross-disciplinary and will involve researchers from networking, systems, and HCI, as well as industry.

Acknowledgments

This article grew out of the 2009 University of Washington/Microsoft Research Summer Institute (<http://www.cs.washington.edu/mssi/2009/index.html>) where researchers from academia and industry explored opportunities in home networking. We thank all participants, as well as our colleagues and collaborators at our respective institutions. The Georgia Tech authors are supported by National Science Foundation awards #CNS-0626281 and #IIS-0904431. 

References

1. ABI Research. *Home Automation and Security Technical Report*. New York, 2009; <http://www.abiresearch.com/>
2. Accenture. *Big Trouble with ‘Trouble Found’ Returns* Technical Report. Dublin, Ireland 2008; http://www.accenture.com/SiteCollectionDocuments/PDF/22701_ReturnsRepairsRvn_v04lr.pdf
3. Aggarwal, B., Bhagwan, R., Das, T., Eswaran, S., Padmanabhan, V., and Voelker, G. NetPrints: Diagnosing home network misconfigurations using shared knowledge. In *Proceedings of the Sixth USENIX Symposium on Networked Systems Design and Implementation* (Boston, Apr. 22–24). USENIX Association, Boston, 2009, 349–364.
4. Blumenthal, M. and Clark, D.D. Rethinking the design of the Internet: The end-to-end argument versus the brave new world. *ACM Transactions on Internet Technology* 1, 1 (Aug. 2001), 70–109.
5. Calvert, K.C., Edwards, W.K., and Grinter, R.E. Moving towards the middle: The case against the end-to-end argument in home networking. In *Proceedings of ACM Workshop on Hot Topics in Networks* (Atlanta, Nov. 14–15). ACM Press, New York, 2007.
6. Edwards, W.K. and Grinter, R.E. At home with ubiquitous computing: Seven challenges. In *Proceedings of the Third International Conference on Ubiquitous Computing* (Atlanta, Sept. 30–Oct. 2). Springer-Verlag, Heidelberg, 2001, 256–272.
7. Glerum, K., Kinshumann, K., Greenberg, S., Aul, G., Orgovan, V., Nichols, G., Grant, D., Lohle, G., and Hunt, G. Debugging in the (very) large: Ten years of implementation and experience. In *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles* (Big Sky, MT, Oct. 11–14). ACM Press, New York, 103–116.
8. Grinter, R.E., Edwards, W.K., Chetty, M., Poole, E.S., Sung, J.-Y., Yang, J., Crabtree, A., Tolmie, P., Rodden, T., Greenhalgh, C., and Benford, S. The ins and outs of home networking: The case for useful and usable domestic networking. *ACM Transactions on Computer-Human Interaction* 16, 2 (June 2009), 1–28.
9. Grinter, R.E., Edwards, W.K., Newman, M., and

- Ducheneaut, N. The work to make a home network work. In *Proceedings of Ninth European Conference on Computer-Supported Cooperative Work* (Paris, Sept. 18–22). Springer-Verlag, Heidelberg, 2005, 469–488.
10. Intel. *Intel Remote PC Assist Technology (Intel RPAT) Software*; <http://www.intel.com/support/services/rpat/>
11. Kandula, S., Mahajan, R., Verkaik, P., Agarwal, S., Padhye, J., and Bahl, P. Heideid diagnosis in enterprise networks. In *Proceedings of the ACM SIGCOMM Conference on Data Communication* (Barcelona, Aug. 17–21). ACM Press, New York, 2009, 243–254.
12. Kempton, W. Two theories of home heat control. Cultural models in language and thought. *Cognitive Science* 10, 1 (1986), 75–90.
13. Laszlo, J. *Home Networking: Seizing Near-Term Opportunities to Extend Connectivity to Every Room Technical Report*. Jupiter Research, 2002.
14. Liu, Z., Lee, B., Kandula, S., and Mahajan, R. NetClinic: Interactive visualization to enhance automated fault diagnosis in enterprise networks. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology* (Salt Lake City, Oct. 24–29). IEEE, New York, 2010, 131–138.
15. Norman, D.A. *The Design of Everyday Things*. Doubleday, New York, 1990.
16. O’Brien, J., Rodden, T., Rouncefield, M., and Hughes, J. At home with the technology: An ethnographic study of a set-top-box trial. *ACM Transactions on Computer-Human Interaction* 6, 3 (Sept. 1999), 282–308.
17. Pew Internet and American Life Project. *Home Broadband Adoption*. Washington, D.C., 2009; <http://www.pewinternet.org/Reports/2009/10-Home-Broadband-Adoption-2009.aspx>
18. Pew Internet and American Life Project. *When Technology Fails*. Washington, D.C., 2008; <http://www.pewinternet.org/Reports/2008/When-Technology-Fails.aspx>
19. Pogue, D. Hot-spot shortcut in the weeds. *New York Times* (Apr. 7, 2010); <http://www.nytimes.com/2010/04/08/technology/personaltech/08pogue.html>
20. Poole, E.S., Chetty, M., Morgan, T., Grinter, R.E., and Edwards, W.K. Computer help at home: Methods and motivations for informal technical support. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems* (Boston, Apr. 4–9). ACM Press, New York, 2009, 739–748.
21. Poole, E.S., Edwards, W.K., and Jarvis, L. The home network as a socio-technical system: Understanding the challenges of remote home network help. *Journal of Computer-Supported Cooperative Work* 18, 2–3 (2009), 277–299.
22. Scherf, K. *Parks Associates Panel on Home Networking* at the Consumer Electronics Association Conference (San Francisco, Oct. 14–16, 2002).
23. Shehan, E. and Edwards, W.K. Home networking and HCI: What hath God wrought? In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (San Jose, CA, Apr. 28–May 3). ACM Press, New York, 2007, 547–556.
24. Steinberg, D. and Cheshire, S. *Zero Configuration Networking: The Definitive Guide*. O’Reilly Media, Inc., Sebastopol, CA, 2005.
25. Tolmie, P., Pycock, J., Diggins, T., MacLean, A., and Karsenty, A. Unremarkable computing. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (Minneapolis, Apr. 20–25). ACM Press, New York, 399–406.
26. Venkatesh, A. Computers and other interactive technologies for the home. *Commun. ACM* 39, 12 (Dec.1996), 47–54.

W. Keith Edwards (keith@cc.gatech.edu) is an associate professor in the School of Interactive Computing of the Georgia Institute of Technology, Atlanta.

Rebecca E. Grinter (beki@cc.gatech.edu) is an associate professor in the School of Interactive Computing of the Georgia Institute of Technology, Atlanta.

Ratul Mahajan (ratul@microsoft.com) is a researcher in Microsoft Research, Redmond, WA, and an affiliate professor in the Department of Computer Science & Engineering of the University of Washington, Seattle.

David Wetherall (djw@uw.edu) is an associate professor in the Department of Computer Science & Engineering of the University of Washington, Seattle.