Moving Toward the Middle: The Case Against the End-to-End Argument in Home Networking

Kenneth L. Calvert* W. Keith Edwards† Rebecca E. Grinter†
*Lab for Advanced Networking, University of Kentucky
†School of Interactive Computing, Georgia Tech
calvert@netlab.uky.edu, keith@cc.gatech.edu, beki@cc.gatech.edu

I. Introduction

Home broadband adoption is growing rapidly in much of the developed world [23], leading to increasing use of networking *inside* homes to allow sharing of devices, content and of the Internet connection itself by multiple hosts inside the home. The increasing penetration of home networking enables new products and applications in health care, entertainment and security, as well as in areas currently unforeseen.

Unfortunately, all is *not* well in the connected home, and the set of problems facing users of home networking is becoming a key impediment to the promised benefits of home networking. Consumer statistics illustrate the problem: In 2003, home networking gear was the *most returned item* at "big box" electronics stores in the US [17]. In 2006, a quarter of wireless access points purchased by consumers were returned [12]—not because the devices "broke", but because their users were unable to properly set up and integrate them into their home networks; complexity is reliably cited as the key impediment to home networking [11]. Even if deployed successfully, home nets are often misconfigured, posing a threat to the larger Internet as evidenced by the existence of large-scale "botnets."

Why should these problems be the domain of the networking research community? Isn't this a simply an "implementation problem" that vendors and ISPs will address? We claim that solving the problems associated with home networking is not simply a matter of building a better home router, providing more clear documentation, or better technical support. Rather, many of these problems arise from a mismatch between certain aspects of the present Internet architecture and the characteristics of the home environment. In particular, the end-to-end principle, which has shaped many of the design choices of the Internet [16], is a contributing factor to the difficulties seen in the home (a point that has been raised by others [18], perhaps most notably by Blumenthal and Clark [3]). We argue that a more comprehensive *system* solution is required.

In this paper we examine the problems that arise from "bringing the Internet home"—applying the same protocols and architectural principles designed for the Internet-at-large in the home environment. This examination is rooted in a series of studies that have explored the problems of networking from a human perspective. We tie these end-user visible problems to the architectural choices that result from the end-to-end argument, and present a new architecture for the home network that elevates ease of installation and use, evolution, and troubleshooting to the same level of importance as the application neutrality and core network simplicity goals dictated by the end-to-end argument.

The rest of this paper is organized as follows. Section II considers why the Internet architecture was brought into the home initially, as well as why it might be inappropriate for the home setting, based on user studies of the travails of networking in the home. Section III presents a set of design goals for any home network architecture. We then present a prototype design and argue that it satisfies the goals. We discuss related work in Section V, and conclude with some thoughts on the implications of our work.

II. The Internet at Home

In this section we consider the use of the current Internet protocols and architecture in the home network environment.

A. Why?

It is understandable why the Internet Architecture "came home" in the 1990's. The basic assumptions of the Internet had by then already proven themselves across millions of nodes, and so extending the architecture into the home was a natural next step. Home devices would fully support the end-to-end nature of the Internet, acting as full-fledged TCP/IP endpoints (modulo complications such as NAT boxes, of course); this approach would allow complete compatibility with the growing range of services on the network.

Many of these same assumptions hold today, of course, which means that the use of Internet protocols and architectural models in the home is a natural, perhaps inevitable, step, for several reasons. First, rest of the world, including the outside applications that residents want to use, uses these protocols, and will for the foreseeable future. Second, the widespread use of the architecture means that inexpensive, embeddable protocol implementations are now available. Third, the "dumb network, smart endpoints" design philosophy of the Internet is supposed to promote simple, robust networks, allowing deployment of new applications easily.

B. Why Not?

Despite the obvious appeal of simply extending the Internet Architecture into the home, there are a number of problems with this approach. Earlier work by our (latter two

authors') group detailed a series of user studies designed to assess householders' conceptions of—and experiences with—networking in the home [10], [4]. These studies used a range of data collection and analysis methods to yield an in-depth exploration of the home networking practices of 42 users over 18 households, and revealed a number of commonly experienced problems. Here, we link these user-visible issues back to their technical roots in the end-to-end Internet architecture.

1) Provisioning: Perhaps unsurprisingly, users experience deep problems provisioning clients for the network. This provisioning includes not only network-layer settings (IP addresses, subnet masks, default routers) and link-layer settings (SSID, WEP keys), but also a large and ever-expanding array of application-layer settings (default printer shares, locations of networked media adapters, file shares, application-specific firewall and NAT settings, and so forth).

We argue that the provisioning problem, taken as a whole, is inherent in the Internet architecture. The end-to-end assumptions made by the Internet architecture rely on less intelligence in the network, and more intelligence at the edges of the network. Unlike a telephone—which can simply be plugged into a wall socket and will work—these smart edge nodes must be configured in order to work correctly on the network; of course, where there is the necessity of configuration, there is the possibility of *misconfiguration*.

To a partial degree, some of these problems can be mitigated through technologies such as DHCP, which remove aspects of provisioning from the hands of users and place them in a per-network authority that can supply the "correct" details for the home network. However, DHCP only solves a small portion of the overall problem, and for most aspects of provisioning there is no equivalent technology. The result is a large and growing array of "bandage" technologies, each intended to address a small slice of the provisioning problem (such as means to provision wireless security keys and SSIDs [7] [2]), rather than provide a holistic solution.

2) Topological Complexity: In theory, the conceptually simple network core argued for by the end-to-end approach should mean that the networking infrastructure itself is relatively pain-free in the home. In practice, however, network topology—physical and logical—causes many problems for users. Creating a functioning network within the home increasingly involves a wide array of infrastructure equipment (switches, hubs, internal access points, powerline bridges), all of which increase the apparent complexity of the system for users.

Part of the problem is that, despite the neutrality to application semantics argued for by the end-to-end model, the topology of the network *does* matter to the applications that use it, and to the hosts that are on it. Topology forces upon users a whole range of new complexities with which they must cope. For example, they may need to understand the implications of having multiple DHCP servers on their network (provided by separate routers or access points, for instance). They may have to understand why multicast traffic (important for discovery protocols) does not cross subnet boundaries—and indeed, even what subnets are and why the addition of a piece of infrastructure equipment may "break" an application like music sharing [21]. They have to understand the difference between the "inside" and "outside" of the home network, and why this logical boundary may not correspond exactly to the physical interconnection of infrastructure devices (open access points can allow "outside" machines to join the home network; without a NAT or firewall devices may be effectively "outside" the home logically although not physically). In short, far from providing a simple, reliable, semantically-neutral service, the network at home is distinctly visible and problematic to users: increasing topological complexity breaks applications, complicates details such as address assignment, and requires extra security awareness and configuration.

3) Troubleshooting: Unlike the technically sophisticated, managed environments in which the Internet was born, the home is a place of relatively unsophisticated users (from a networking perspective), who are little interested in network management as an end in itself. Networking problems are a bane to these "reluctant administrators," who are often confused about where to even start troubleshooting. If a laptop on the home network cannot access the web, for example, it may be hard to determine whether the problem resides with the laptop, the wireless access point, the home router, the ISP, or the website. Certainly the network itself provides little help in diagnosing problems in the application terms that users understand. This is at least partly a consequence of the network core's application-neutrality—indeed, the network provides very little that could be helpful in troubleshooting, even for the most common/standard transport protocols (TCP). As others have observed [5], there is a need for monitoring and management aid that extends beyond the network endpoints

Troubleshooting difficulties are often compounded by the other issues discussed in this section. Increased topological complexity frustrates diagnosis for both householders and for remote service providers, who have to help without knowing the intricate details of the individual network. Topological complexity also makes it difficult to collect data for diagnosis, as no one node may have a complete picture of traffic on the home network. Another compounding factor is that the network infrastructure is more or less invisible to its users. For one thing, the logical configuration of the network cannot be determined by simply looking at it—even for an expert. Worse, the *physical* infrastructure may be functionally invisible: infrastructure devices were often hidden under sofas and in closets in our studies. Many householders in our studies, for example, were entirely unaware of many of the infrastructure devices on their own networks.

4) Security: As has been widely observed, the Internet architecture was designed when it was reasonable to assume some level of trustworthiness at endpoints. While this has not been the case for some time now, the security solutions provided by current home networking technologies are not well suited to the needs or abilities of their users. Even basic mechanisms—such as setting up WEP—are often not enabled in the home, not out of a lack of interest but because of complexity. Security problems, however, go far beyond simply the failings of current wireless technology. The de facto setup of most home networks involves a NAT device that defines the border of the network, separating devices "inside" the network (assumed to be trustworthy) from devices "outside" (against

which the home machines must be protected).

This conflation of topology with trust breaks down in a number of ways. First, some applications require opening home machines, which are generally weakly defended, to public Internet access; not only is this strategy risky, it is also hard for householders to implement, as it involves NAT forwarding, firewall configuration, and so on. Second, this "crunchy on the outside, soft on the inside" approach means that hosts on the home network—whether intentionally, as in the case of guest access, or accidentally, as in the case of neighborhood machines associating with a home's access point—generally have unfettered access to services and data residing on the home network. There is no separation between *access to the network* and *access to the services on that network*. This de facto policy mechanism is a poor fit for householders' needs, as well as their social practices. Support for more fine-grained policy is needed; unfortunately the Internet architecture lacks the mechanisms required to support such policies, and instead overloads mechanisms (addresses and ports) intended for other purposes. Many implementations even require that policies be *specified* in terms of such concepts.

5) Composition: A final issue concerns the difficulty of composing together the increasingly rich variety of networked devices available to consumers. The end-to-end approach has motivated the creation of a diverse range of smart endpoints for the home, all of which can be deployed without changes to the network core. However, the promise of such devices is often frustratingly out-of-reach because of the inability of users to get them to work together.

The ene-to-end principle suggests endpoint-mediated composition, which in turn requires endpoints to have knowledge of each other. Unfortunately this often means software installations, upgrades, and management by end users are necessary to achieve the compatibility necessary for composition. Furthermore, the introduction of any new type of device onto the network may necessitate upgrade of all of the *existing* nodes in order to allow them to use the new device [8]. The end result is isolated "islands of interoperability" in the home. Such interoperability problems are a presently significant barrier to the adoption of new applications and technologies in the home. Centralizing some of the required functionality in the home network "core" could provide a way around this problem.

C. Moving Forward

One could argue that the best way forward is to keep the current Internet architecture in the home and "patch" it in much the same way that DHCP adds "side protocols" that remove some of the problems we have observed. We argue that while a "bandage" strategy may work for a select subset of the problems enumerated above, it cannot address the broader challenges of usable home networking. Instead, solving the home networking conundrum calls for a complete rethinking of the architecture and in particular the assignment of functions to components. The main argument for this approach is the observation that a networking architecture has many goals, and the relative priority of those goals in the home is quite different from the context in which the Internet originally evolved. In recent years, the same observation has been made for other environments, including sensor networks [15] and delay-tolerant networks for space exploration [9]. We propose that, like these other sorts of networks, the home network should represent a distinct sort of "edge" network: potentially using a different architecture internally, yet able to connect to the larger Internet through specialized gateway devices.

The next section lays out both requirements and challenges for a new home network architecture. Following this we discuss a prototype design for such a home network that we are creating.

III. Requirements for a Solution

In this section we present a set of requirements for an architecture designed specifically for the home.

- **Self-configuring, self-administering.** As noted earlier, provisioning and configuration are probably the biggest challenges for home networks. To the extent possible, these responsibilities must be shifted away from humans, to prevent the possibility of mis-configuration. Instead, human action should be confined to only those tasks that the system cannot infer itself. These tasks will largely be centered around setting *policy*—which machines are a part of the network and which are not, which devices are allowed to communicate with the outside world, etc.
- **Secure by default.** Products sold by individual vendors generally take a liberal security stance: anything not expressly forbidden through user intervention is permitted. This is understandable from the point of view of making products easy to configure and use (and therefore easy to sell) but in the long run it leads to a less usable Internet for *everyone*. Therefore the home network must be "automatically" secure, without the need for human involvement—it should be very difficult or impossible to set up the network in such a way that it can be used without being secured. This requires a more conservative security stance. In particular, the act of joining the network must be restricted to explicitly-authorized clients, and communication among clients must be explicitly enabled.
- **Explicit user interface.** As observed in the previous section, the virtual "invisibility" of current network infrastructure makes troubleshooting difficult. When there *is* an explicit "manifestation" with which users can interact, it typically takes the form of web servers scattered across various devices (which may be unreachable in case of trouble) or flashing lights on front panels. The network must provide a single, well-known, flexible device through which users can interact with and control all aspects of their network.
- **Compatible with existing external TCP/IP-based applications.** Clearly any new architecture for the home must allow the use of the existing services (web and mail servers, games, and so forth) on the current Internet. Note however that compatibility with existing devices and services intended for use *inside* the home is *not* a requirement of ours.

- **Application-independence.** An extremely important feature of IP is its obliviousness to the characteristics of any particular application. This "Prime Directive" of networking is the core of the end-to-end principle, and the feature that makes the network robust and evolvable, enabling it to support applications that were undreamed-of when it was invented. The home network must also have this characteristic, especially since the home is likely to be the focal point of more and more networked applications. Thus—title of this paper notwithstanding—rather than throwing out the end-to-end argument entirely, we take a "strict constructionist" view.
- **Support for composition.** Ultimately, users deploy home networks because of the value that those networks promise to provide; often this value is in the ability to interconnect devices within the home to share media and data. As observed in the previous section, a strict end-to-end approach fundamentally makes this problem harder for *both* application equipment vendors and users, because devices cannot depend on on the network infrastructure to help; they must agree among themselves, *in advance*, on ways to compose their respective applications. A better approach is for the home network service to provide generic facilities, on top of which tools for easier device and service composition can be created.

IV. A "Smart Middle" Design

In this section, we present a prototype architecture for the home that departs from the traditional "dumb middle, smart ends" approach widely credited for the Internet's success.

A. Required Functions

Based on the foregoing observations, we conclude that the home network must provide the following five functions:

1) **Packet transport**, both among endpoints "inside" the home and between inside and outside and endpoints.
2) **Status monitoring**, to assist in troubleshooting. Ideally, the network "knows" what every connected device is, where it is, and with what other devices it has communicated (or not) recently.
3) **Policy enforcement and solicitation**. The network must enforce the home's policies regarding allowed and prohibited communication. Because trivial default policies are generally either not safe ("everything is allowed") or not useful ("nothing is allowed"), aid must be provided for setting sensible, home-specific policies for both intra- and extra-home communication.
4) **Brokering**, to enable devices to discover others of interest on the network (modulo policy constraints). When a new audio output device connects to the network, for example, it should learn about all audio sources (HDTV receiver, CD player, DVD player) with which it is compatible.
5) **User Interface**. As described in the previous section, the network system must have a simple, explicit "manifestation", including a means of controlling and troubleshooting the network.

The network system must provide these functions in a manner that is consistent with the "Prime Directive", i.e. without constraining future applications. Also, the design must be compatible with a wide variety of link-level technologies, from wired Ethernet to WiFi to optical fiber.

B. Distribution of Function

The Internet infrastructure provides a simple, best-effort service, with higher-level functionality delegated to the edges. The principle behind this design is that building application-specific or advanced functionality into the infrastructure can sometimes get in the way of providing the basic service needed by future applications, or of scalability. A loose construction of this principle might suggest that each home device should take responsibility for finding others with which it needs to interact, for maintaining and enforcing its own (and others') policies governing allowed communication, for providing its own user interface, and for assisting the user with troubleshooting. These functions can be much more efficiently provided with some kind of shared infrastructure. The principle again suggests that maximum flexibility results from having this infrastructure in the "ends". Unfortunately, as we have noted above, flexibility is actually a hindrance when it comes to inferring network topology and isolating trouble. Indeed, it has been suggested that the "dumb middle, smart ends" principle is behind many difficulties with network management in general [5].

The status quo in home networking—characterized by a $79 box providing DHCP, NAT, port forwarding, and possibly DNS service, in addition to basic connectivity—represents an attempt (even a valiant one) to provide infrastructure for some or all of our five functions in a mass market "end" device. The problem is that the architecture does not adequately constrain the space of possible topologies and configurations, so there are few true invariants that such a device can count on. In particular, the one function it seems very difficult to provide is assistance with troubleshooting. Here, the inability to assume anything about topology makes it difficult for such devices to provide meaningful help with problem diagnosis.

For these reasons, our proposed home network design takes a "smart middle, smart ends" approach. The required functions are provided by a central component we call the *portal*, which provides basic connectivity, controlling and mediating all communication among devices inside the home, as well as between the inside and outside of the home, through its "interconnect" component (Figure 1). It stores and enforces all policy relevant to the network, including which devices inside the home are allowed to receive communications from outside, and under what circumstances. It maintains a database of device and location information, and provides a brokering service to devices that need to interact with each other. It gives the network an explicit "presence" in the home through a simple user interface. Finally, but perhaps most importantly, it monitors the status of all devices connected to the network, and can assist with troubleshooting. In addition to its centralized, integrated infrastructure, the architecture has two other noteworthy features.

First, an *explicit introduction step* is required for a user to "introduce" a device to the network. During this step, the portal and the device communicate over a physically secure channel, to accomplish the following:

- Exchange encryption keys and lists of supported cipher suites. Maintains the invariant "every connected device has a secure channel to the portal."
- Device downloads a list of exported functions, uploads a list of existing functions on the network.
- Portal learns other attributes from device/user: name, mobile/stationary status, location in home, etc.
- The device informs the portal of any *policies* relevant to its operation that need to be specified; the portal prompts the user (e.g., via a series of questions) to configure them.
- The device downloads into the portal any code needed for mediating its communication with the outside world.

The last item indicates the other salient feature of the portal: *extensibility*. The portal exposes an API that allows it to provide assistance—for example, application-level gateway functionality—for the device's communication needs. This API provides a narrow interface, giving access to needed capabilities that are nevertheless limited, to prevent mischief.Example uses of this API include: conversion to allow communication with legacy devices, and delegation of TCP endpoint functionality so a simple device can communicate with a different protocol on the internal network while still using TCP to communicate with an outside peer.

Figure 1 shows our design, which uses commodity components. The interconnect is provided by a managed gigabit Ethernet switch; we are experimenting with the dynamic use of VLANs to control connectivity among devices. The "user interface" device, envisioned as a kind of universal remote control, is based on a Nokia 700 handheld, communicating wirelessly with the controller. The design of a general, powerful, API for extensibility is part of our ongoing research. As in an earlier prototype [24], the introduction step takes place over a dedicated, secure channel.

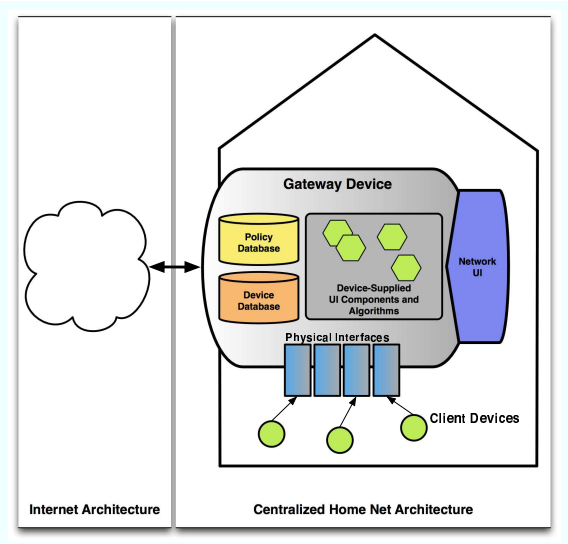

Fig. 1. Portal components

C. Discussion

Here we consider some advantages and challenges of the design. The centralized approach changes device-to-device discovery and communication from an $O(N^2)$ to an $O(N)$ problem. However, the portal represents a single point of failure for the home network—if the device breaks, not only is connectivity to the Internet severed, but communication between devices *within* the home network is also prevented. While such a single point of failure can be seen as a disadvantage of our approach, it also offers some advantages. Unlike in current networks, where users often have no good way of isolating points of failure, the portal-mediated architecture strictly limits the number of failure modes. If devices cannot communicate with each other, either (i) one or both of them has failed, (ii) some channel has failed, or (iii) the portal itself has failed. If the portal fails, functionality will be unavailable and the culprit will be obvious; otherwise, its monitoring capability can provide assistance in identifying the failed component. Of course, portal failure should be rare; to ease replacement when necessary, there must be some means to securely extract and transfer policy settings and other configuration information to the replacement device.

A second concern is performance. Routing all traffic through the central interconnect may incur a performance hit, especially with shared-medium wireless substrates. While bandwidth requirements for home networks today are relatively modest, the increasing penetration of HDTV may press the issue fairly soon. However, we believe that the size of the home network mass market will provide strong incentives to solve this problem.

Third, the central portal must have all of the physical interfaces that are (or will be) necessary in the home. This requirement raises issues of forward compatibility: as new interface types are defined, the portal must be easily expandable or replacable as necessary. This incurs an expense, as does the need to implement the wired home run infrastructure.

Our proposal introduces an architectural discontinuity at the home boundary. The portal must be responsible for ameliorating the effects of that discontinuity. A programmable API that allows device-specific mediation is a key component of the solution, and another subject of ongoing research.

V. Related Work

Despite widespread difficulties with home networking, surprisingly little research has focused on how householders experience networking problems. The previously cited studies by our group [10], [4], as well as a handful of others (such as [20]) represent the only systematic investigations of the user experience of home networking to our knowledge.

A number of efforts have focused on making networking easier for end-users. With regard to provisioning, DHCP is a core part of most home networks, and relieves householders of a number of aspects of manual configuration. IPv6's stateless configuration mechanisms also provide automatic address assignment. Both, however, are limited to provision of basic network-layer information. Other projects have addressed selected other aspects of provisioning, such as SSID and WPA

keys [7] and 802.1x certificates [2]. Although all of these increase convenience for aspects of the end-user configuration problem, in total they represent a hodge-podge of solutions, each addressing a small portion of the overall problem.

A range of technologies layered atop TCP/IP aim to provide better composability. Two key examples are Universal Plug and Play [14] and Jini [22], which provide service architectures allowing devices on the network to discover and interoperate with one another. Both, however, require that clients have knowledge of the device types with which they will communicate. Further, both focus largely at the application layer, not on the lower-layer complexity issues faced by users.

Zeroconf (also known as Bonjour) provides service discovery based on multicast DNS. Although it has been used in a range of applications, even this technology can cause problems for users since multicast traffic may not cross subnet boundaries. The rapidly growing complexity of the home network frustrates simple solutions such as this; in our studies, for example, users reported confusion around understanding the limits of discoverability, an issue directly caused by the increasing topological complexity of the home network.

The Open Services Gateway Initiative (OSGi) is a modular and easily adaptable service architecture for the home [1]. In OSGi, a residential gateway is provisioned with code in the form of "service bundles" that extend its behavior. While OSGi's centralized, extensible gateway is similar to our portal, OSGi requires an already functioning IP-based network in the home in order to work; it does not address issues of management and complexity of the IP network itself.

Network management has a long history in the networking community, albeit mainly in the enterprise context. Systems such as HP OpenView [6] provide powerful management consoles built on top of SNMP [19]. These sorts of technologies are inappropriate for home users, who have neither the technical sophistication nor motivation to deal with this degree of management complexity. The work of Clark et al on a "knowledge plane" for the network [5] is an attempt to deal with network management issues that result from the "dumb middle" approach, by adding infrastructure to monitor and assist with troubleshooting. While the design of such infrastructure is far from obvious for the global Internet, the reduced scale and performance requirements of the home network make a simplified, parsimonious design possible.

Very recently, a number of systems have begun to appear that provide a home-centric view of the network. Pure Networks Network Magic tool and the network monitor in Windows Vista provide simple visualizations of the home network [13] as well as troubleshooting guidance. While certainly helpful, these tools are necessarily limited by the constraints imposed by the complexity of the Internet-based home network architecture. For example, these tools cannot collect data on non-local links on the home network.

VI. Conclusions

We have argued that the unique characteristics of the home environment require a refactoring of network functionality. Our proposed design extrapolates the de facto standard common home network architecture, which features a commodity multifunction box as the central component, to provide crucial functions missing from the Internet architecture. We believe that this shift away from a pure end-to-end approach, toward one in which the network core gains much greater functionality, can address many of the empirically observed problems with networking in the home.

Acknowledgement. This work was supported by the National Science Foundation (CNS-0625802 and CNS-0626281).

References

[1] Open Service Gateway Initiative Alliance. Osgi–the dynamic module system for java.

[2] Dirk Balfanz, Glenn Durfee, Rebecca E. Grinter, Diana Smetters, and Paul Stewart. Network-in-a-box: How to set up a secure wireless network in under a minute. In *Proceedings of the USENIX Security Symposium 2004, San Diego, USA*, pages 207–222, August 2004.

[3] M. Blumenthal and D. D. Clark. Rethinking the design of the internet: The end-to-end arguments vs. the brave new world. *ACM Transactions on Internet Technology*, 1(1):70–109, August 2001.

[4] Marshini Chetty, Ja-Young Sung, and Rebecca E. Grinter. How smart homes learn: The evolution of the networked home and household. In *Proceedings of the Ninth International Conference on Ubiquitous Computing (Ubicomp), Innsbruck, Austria*, September 2007.

[5] D. Clark, C. Partridge, J. Ramming, and J. Wroclawski. A knowledge plane for the internet (position paper). In *Proceedings of ACM SIGCOMM 2003, Karlsruhe, Germany*, pages 3–10, August 2003.

[6] Hewlett-Packard Corporation. Hp openview management software.

[7] Linksys Corporation. Secureeasysetup.

[8] W. Keith Edwards and Rebecca Grinter. At home with ubiquitous computing: Seven challenges. In *Proceedings of the Conference on Ubiquitous Computing (Ubicomp 2001), Atlanta, GA*, September 2001.

[9] Stephen Farrell, Vinny Cahill, Dermot Geraghty, Ivor Humphreys, and Paul McDonald. When tcp breaks: Delay- and disruption-tolerant networking. *IEEE Internet Computing*, July/August 2006.

[10] Rebecca E. Grinter, W. Keith Edwards, Mark W. Newman, and Nicolas Ducheneaut. The work to make the home network work. In *Proceedings of the 9th European Conference on Computer Supported Cooperative Work (ECSCW '05), Paris, France*, pages 469–488, September 2005.

[11] J. Laszlo. *Home Networking: Seizing Near-Term Opportunities to Extend Connectivity to Every Room*. Jupiter Research, July 2002.

[12] R. MacMillan. Wireless networking baffles some customers. Reuters News Report, March 2006.

[13] Pure Networks. Network magic.

[14] Universal Plug and Play Forum. *www.upnp.org*.

[15] Kay Romer and Mattern Friedemann. The design space of wireless sensor networks. *IEEE Wireless Communications*, 11(6):54–61, 2004.

[16] J. Saltzer, D. Reed, and D.D. Clark. End-to-end arguments in system design. *ACM Transactions on Computing Systems*, 2(3):277–288, November 1984.

[17] K. Scherf. Panel on home networking. Consumer Electronics Associate (CEA) Conference, 2002.

[18] Erika Shehan and W. Keith Edwards. Home networking and HCI: What hath god wrought? In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'07), San Jose, USA*, April 2007.

[19] W. Stallings. Addison-Wesley, 1998.

[20] Peter Tolmie, Andy Crabtree, Tom Rodden, Chris Greenhalgh, and Steve Benford. Making the home network at home: Digital housekeeping. In *Proceedings of the European Conference on Computer-Supported Cooperative Work*, 2007.

[21] A. Voida, R. Grinter, N. Ducheneaut, W. K. Edwards, and M. Newman. Listening in: Practices surrounding itunes music sharing. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI05), Portland, OR*, April 2005.

[22] Jim Waldo. The jini architecture for network-centric computing. *Communications of the ACM*, pages 76–82, July 1999.

[23] Richard Wray. Broadband spreads across globe. The Guardian (UK Newspaper), June 2007.

[24] Jeonghwa Yang and W. Keith Edwards. ICEbox: Toward easy-to-use home networking. In *Proceedings of IFIP Conference on Human-Computer Interaction, Rio de Janeiro, Brazil*, September 2007.