



Intel[®] Technology Journal

Interoperable Home Infrastructure

**Meeting the Demands of the
Digital Home with High-Speed
Multi-Hop Wireless Networks**

Meeting the Demands of the Digital Home with High-Speed Multi-Hop Wireless Networks

Lakshman Krishnamurthy, Network Architecture Lab, Intel Corporation

Steven Conner, Network Architecture Lab, Intel Corporation

Mark Yarvis, Network Architecture Lab, Intel Corporation

Jasmeet Chhabra, Network Architecture Lab, Intel Corporation

Carl Ellison, Network Architecture Lab, Intel Corporation

Chuck Brabenac, Communications and Interconnect Technology Lab, Intel Corporation

Ernest Tsui, Communications and Interconnect Technology Lab, Intel Corporation

ABSTRACT

In the near future, homes will be equipped with wireless networks that bridge data and consumer electronics networks, interconnecting desktop PCs, mobile laptops and handhelds, High-Definition TVs (HDTVs), DVD players, camcorders, and other multimedia devices. This environment introduces new wireless network requirements, including high and dependable bandwidth, low latency, and coverage throughout the home. Multi-hop wireless technology offers unique benefits for creating a high-speed, robust home wireless network. However, to support these demanding usage models, significant wireless networking innovations are required across the physical, MAC, and routing layers, and solutions need to be found for higher level issues such as Quality of Service (QoS) guarantees, device discovery, and security. In addition, user acceptance of multi-hop wireless networks will require ease of installation. Intel R&D is currently researching self-organizing multi-hop wireless networks for home environments. This paper introduces the technologies and tradeoffs needed to create a multi-hop wireless home network, identifying benefits and limitations. In particular, we describe usage scenarios and assumptions that drive the requirements. Finally, we provide an outline of the key technology problems that must be solved and recommend the necessary next steps to make this vision a reality.

INTRODUCTION

Wireless Local Area Network (WLAN) technologies are beginning to gain a foothold in the home computer market. Wireless networks allow the home user to share data and Internet access without the inconvenience and cost of pulling cables through walls or under floors and without unsightly network jacks. Pulling cables and installing new network jacks are particularly challenging

in existing homes and apartments. In addition, wireless LANs provide the convenience of untethered computing for laptops and handhelds from anywhere in or around a house.

The benefits of wireless need not be limited to computer networking. As the bandwidth of wireless networks increases, audio/video home entertainment will be the next target, replacing device-to-device cabling as well as providing distribution throughout the home. Rather than maintaining separate networks for different types of devices, as is common with wireless LANs and cordless telephones, a unified technology is desirable to reduce the cost and complexity of installation and to allow cross-device communication and functionality. For instance, a consumer should be able to insert a DVD into a player in the living room and watch it on a TV in the bedroom or on a laptop on the back porch. Such a network will need to span the entire home, allowing any two devices to communicate. Figure 1 below shows typical devices in a home: entertainment devices, HDTVs, DVD players, game consoles, PCs, and laptops. A multi-hop network that interconnects these clusters is also shown.



Figure 1: Devices in a typical home clustered in various rooms

In a multi-hop network, a node transmits with low power to reach nearby neighboring nodes (routers), which will forward the data toward the intended destination. Using a multi-hop network provides significant benefits assuming a limited channel capacity. The alternative is a single-hop network, where each device transmits with enough power to be received by any other device in the home. When two devices transmit simultaneously, the resulting channel contention can limit capacity in a high-bandwidth environment. The typical solution to this problem is to divide the channel into subchannels by some combination of frequency, time, or coding. Dividing the channel into N subchannels allows N devices to transmit without contention. However, each of these subchannels will have a capacity of at most $1/N$, which must be sufficient to carry the desired traffic.

Multi-hop networking, on the other hand, increases the aggregate capacity of the network by using lower transmit power to only reach nearby neighboring nodes. This allows channel re-use, thereby improving spatial capacity. By using lower transmit power, devices at different locations can transmit simultaneously without interference. Thus, despite a limitation of N channels, more than N devices can potentially transmit simultaneously without contention.

In addition to preserving spatial capacity, the low-transmission power requirements of multi-hop networks allow them to support higher bandwidth despite Federal Communications Commission (FCC) regulations that limit maximum transmission power. At a given transmission power level, the number of reception errors increases as transmission distance increases, due to a diminishing signal-to-noise ratio. To allow transmission over greater distances, wireless devices use variable forward error correction encoding schemes or step-down to simpler modulation schemes [4]. These schemes result in a decrease in channel capacity as the transmission distance increases (Table 1). Multi-hop networking avoids this problem by transmitting data over several short hops rather than one large hop.

	802.11a	802.11b	802.11g	UWB
5m	54	11	54	660
10m	48	11	54	188
20m	36	11	48	20
30m	24	11	36	5
40m	18	5	24	2
50m	12	5	18	1
60m	9	2	12	0.5

Table 1: Raw channel capacity (in Mbps) for several wireless technologies at various communication ranges. Higher bandwidths are available at shorter distances irrespective of the technology or standard.

Multi-hop communication also provides greater redundancy. When the network is dense, each device can have many neighbors within communication range, potentially creating multiple paths between two communicating devices. In the presence of localized interference or attenuation, such as a person standing in a room, a multi-hop network can route data along an alternate path. In a single-hop network, it is not possible to route around interferences or degradation between two devices.

While multi-hop networking solves many problems, many challenges remain before it can become a reality. In particular, a home multi-hop networking environment must be self-administered and cannot require the user to be technologically savvy. Currently, the installation of even a single-hop wireless LAN is a challenge [8] [10]. Identifying the optimal location for access points, selecting communication channels, setting the transmit power, and enabling security all require a high level of technical expertise and engineering. This paper discusses approaches to these and other issues in the context of multi-hop networking. While we do not offer definitive solutions we do suggest a number of research directions that will help to make high-speed multi-hop wireless networks for the home a reality.

USAGE SCENARIOS AND REQUIREMENTS

The first step in the process of creating deployable multi-hop networks is to identify the requirements of the home network based on network usage scenarios. Home networks may be deployed in a variety of domains:

- a small house, well separated from other houses
- an apartment in an apartment complex
- a large suburban house, with computer-savvy children
- a townhouse in a row of townhouses
- a college dorm or other similar facilities

Each of these domains may be viewed as a collection of interconnected clusters of devices. We need to answer three questions in connection with this network. First, how does a user install a network in a diverse home environment and maintain interoperability? Second, what are the traffic characteristics of devices and applications in this network? Third, how do multiple

networks co-exist in each domain? The following subsections explore each of these questions.

Interoperability Requirements

The interoperability issue is not unique to multi-hop wireless networks. However, any multi-hop solutions must account for the diversity of devices in the home network. In the future, several classes of wireless networking devices will be purchased by consumers for installation in the home. Consumer electronics and computing devices such as DVD players, televisions, remote-control devices, and handheld computers will come with radios pre-installed (e.g., Radio Free Intel [11]).

In the home environment, the wide range of types of equipment means that we cannot use the same radio everywhere. In many cases, wireless interfaces included in home products should be optimized for a particular task, such as short-range wire replacement [1] [2]. But even these interfaces are diverse. For example, using an expensive high-bandwidth radio on a wireless keyboard is wasteful and will result in a needless increase in cost. On the other hand, a high-bandwidth radio might be appropriate for a DVD player as the increased cost will be small relative to the total cost of the device, and it would support the high-bandwidth usage requirements of the player.

In order to support this rich diversity of devices, the multi-hop network must interface with each kind of device.

Installation Requirements

To enable longer-range multi-hop communication between devices distributed throughout the home, we envision the use of specialized low-cost router devices. Such routers might be packaged in compact form-factors that can be conveniently installed by plugging them into power outlets throughout the home. External add-on radios will convert legacy devices such as home appliances and older audio and video equipment into wireless-enabled devices.

For a non-technical home or apartment dweller to install and configure a home network, it is imperative that multi-hop wireless networking not increase the installation complexity of consumer electronics equipment. Given the large variation and unpredictability of Radio Frequency (RF) propagation in different deployment environments [6] and given the lack of technical expertise by typical installers (homeowners), tools to aid in correct deployment will be very important to ensure good connectivity between devices.

Traffic Characterization

Traffic in this network may be generated by a variety of applications ranging from Internet browsing, data backup, and telephony, to entertainment and gaming. These applications generate a range of traffic patterns.

Interactive traffic: PCs, laptops, and handheld devices will require regular Internet access over a home broadband connection. For applications such as Web surfing, digital photos, and e-mail, bottlenecks are likely to remain in the Internet or on the broadband connection. This traffic has more stringent latency requirements but less stringent bandwidth requirements than entertainment applications. Traffic generated by devices such as wireless mice and keyboards require very low latencies, but these devices are likely to require only single-hop communication over a short distance.

Bulk traffic: Applications such as data back-up, network file storage, and printing require higher bandwidth and, unless metered, may even saturate the available bandwidth. Such traffic would require less priority than interactive applications.

Audio applications: Cordless telephones are common today. Extending these devices to support Internet telephony is a logical next step. This class of applications would require low latency and uninterrupted connectivity while roaming throughout the home.

Entertainment-quality traffic: Audio Visual (A/V) home entertainment devices such as High-Definition TVs (HDTVs), DVDs, stereos, camcorders, multi-media PCs, and musical instruments require high bandwidth. Many of the usage scenarios call for communication in close proximities, but the ability of users to store legal digital content in Personal Video Recorders (PVRs) for distribution around the home introduces a traffic pattern beyond device clusters.

This discussion demonstrates the need for different traffic classification and prioritization within the nodes and the network. We have demonstrated the need to support low-latency and high-bandwidth real-time applications. For multi-hop networking to be useful, it must be able to provide sufficient bandwidth in order to differentiate itself from single-hop solutions.

Coexistence

When multiple networks exist within radio range of one another, these networks must be able to coexist. This situation typically occurs in high-density housing (apartments, college dorms, townhouses). The coexistence question is not unique to multi-hop networks [23]; these networks make the problem harder to solve by requiring a solution at every node in the network.

Coexistence imposes specific requirements for channelization, routing, Quality of Service (QoS), and security.

Channelization

One possible approach to coexistence is the *possessive* one: “I use my network for my devices and you use your network for your devices.” Possessiveness may lead to disaster when it comes to channelization. If two neighbors choose channels for their networks independently, nothing will stop them from choosing the same channels and thus interfering with each other [13].

Channel selection must be cooperative and, to minimize the human administrative effort, should be achieved automatically by the network devices without requiring the network owners to meet and make plans. In a dense enough apartment complex, a network owner might not even know which neighbor is running the competing wireless network.




Figure 2: Two neighboring networks where Nodes 1 through 7 belong to one apartment and Nodes A through F belong to the adjacent apartment

Routing

If the two networks in Figure 2 must remain isolated, the dotted links will not be used. As a result, each network will offer lower performance to its owner than if the two networks were to cooperate in routing. For example, consider the network shown in Figure 2 and assume that the wall shown does not attenuate the signal any more than the internal walls in each apartment (not shown in the figure). Also assume that message transmission time is proportional to the square of the distance covered. Using a simple simulation, we get a throughput improvement between nodes {1, 7} and {2, 3, 4, 5, 6} of a factor of 3.6, while end-to-end latency improves by a factor between 1.5 and 1.9. For link A-B, the throughput improves by a factor of 1.5, and the latency improves by a factor of 1.6. While this result is based on a simplistic simulation, it illustrates the fact that cooperation between

coexisting networks could give enough performance improvement to warrant the effort.

Quality of Service

Even an isolated network requires QoS routing to support streaming audio and video [19]. To achieve the desired QoS in a multi-hop network, every node needs to cooperate with its neighbors whether it is owned by the same user or not. Using low-power radios reduces the severity of the problem by limiting the range, and therefore reducing the number of cross-wall radio neighbors. If we allow neighboring networks to cooperate in order to gain a performance advantage, then we must ensure fairness of network usage. This, too, is a QoS problem that depends first on defining “fairness.”

Security

In a multi-hop network, security is required to enable four major protection functions:

End-device and router introduction: When new end-point devices (e.g., DVD players and music jukeboxes) or router nodes are added, their introduction must be authenticated. This essentially determines the notion of who owns a particular device. This problem is not unique to a multi-hop network, but solutions to the introduction problem must work across the network.

User data integrity and secrecy: Link-level encryption has been proposed for the protection of both data and access in single-hop networks [4]. As discussed in a later section, end-to-end encryption is more appropriate for protecting user data than link-level encryption.

Device control and authentication: Commands sent to devices in a network must also be authenticated. In a wireless network, it is particularly important for end-devices to authenticate users before granting access and control permissions. For example, nodes in a neighboring network should not be allowed to control the television next door, nor should they be able to access personal home movies stored on a neighbor’s media server. These issues must be solved in the context of a multi-hop network.

Network authentication and coexistence: Solving the network authentication problem is especially important with respect to the coexistence problem. A hostile neighbor or intruder could introduce inaccurate routing information or inject an unauthorized traffic load. Packets containing routing updates or QoS-protected streams must be authenticated. Implicit authentication by encryption is a poor substitute for real authentication. Moreover, relying on link encryption is a poor choice in multi-hop networks as end-devices lose access to origin authentication information.

SOLUTION SPACE

In the previous section, we identified the requirements for interoperability, installation, supporting home multi-media traffic, and coexistence. This section looks at the solution choices and tradeoffs that meet these requirements.

Interoperability

In the requirements section, we noted that a typical home could have a variety of devices using a mixture of Physical layers (PHY) and Medium Access Control (MAC) layers, not necessarily directly interoperable. In this diverse home environment, we envision a multi-hop network, using dedicated homogenous router devices. These routers use the same PHY/MAC layer for inter-router communication.

Using the same PHY/MAC in the multi-hop backbone provides an opportunity to better control the network, which makes it easier to provide entertainment-quality connectivity throughout the home. However, the multi-hop network must still interoperate with the wide variety of home devices. As shown in Figure 3, the leaf nodes of the multi-hop network must support every possible PHY/MAC standard that may be used in the home.



Figure 3: A multi-hop backbone must interface with multiple PHY/MAC layers in the home

This issue leads to an important cost-complexity tradeoff. Providing every possible PHY/MAC hardware implementation in each router node is simply too expensive. Creating many different kinds of router devices, each with a particular PHY/MAC implementation increases the complexity of the installation and limits support for mobility. An alternative is to create one type of device (or a small number of them) combining a limited number of PHY/MAC choices based on the likely set of home devices. The cost of this approach needs to be traded off

with a third approach that makes use of reconfigurable radios.

Reconfigurable radio, or Software-Defined Radio (SDR) technology, allows a single piece of silicon to be reconfigured to implement many different PHYs and MACs. Such a flexible radio can allow interoperability with a larger set of end-point devices at a lower cost than including multiple radios in the same device. More information on reconfigurable radios and network configuration protocols needed to support self-configuration in the network may be found in Appendix A; these issues also apply to single-hop networks.

Multi-Hop Network Installation

Using software-defined reconfigurable radios will address the issues of legacy equipment and non-interoperable wireless standards to some degree. However, as described in the requirements section, multi-hop wireless networks must also support ease of installation and placement of nodes in a multi-hop network.

A rule of thumb based on typical deployment scenarios could be supplied to users as a starting point. An example of such a rule could be that wireless routers should typically be deployed every 10 feet. However, in a real home, Radio Frequency (RF) shadows are likely to exist due to home furnishings, household items, people, metal, and other attenuators built into the building structure, thereby greatly limiting the usefulness of such rules of thumb.

Given such unpredictability, users will need help deciding where specifically to deploy their nodes. One possibility is to provide a feedback mechanism, perhaps through Light-Emitting Diodes (LEDs), indicating the signal strength on each node. Such an approach would be particularly helpful for one-hop networks, where the user must simply make sure each device is close enough to another device. However, in the multi-hop router case, each router must be strategically placed to provide sufficient connectivity among multiple nodes (often in more than one direction). One technique to verify signal strength between specific pairs of nodes is to install a switch on each node allowing the user to select a single pair of nodes at a time to verify their connectivity. Using this technique to deploy even a few nodes in a home may become tedious.

An alternate technique would be to deploy the initial network using a rule of thumb, and then to connect a PC or other specialized network monitoring device to the network, which would collect signal strength and connectivity statistics from each node in the network and display a simplified summary of the results to the user.

Any nodes that are not discovered at the network monitoring location indicate a network partition. This information would let the user know that some routers must be moved, or that more routers must be installed in the area between the detected and missing nodes. Nodes that are expected to be in communication with one another, but which have low inter-communication signal strength, should be supplemented by placing a router between them. Such a network monitoring device would make it easy for users to gain a global view of how well they deployed their network.

SUPPORTING DIGITAL HOME TRAFFIC

Beyond the interoperability and installation issues, the Quality of Service (QoS) need of traffic described in the requirements section must be met. Towards this end, we look at the alternatives and tradeoffs in the area of routing, QoS support and channelization.

Routing and QoS Support

The topology of a multi-hop wireless network is a graph with an edge between each pair of nodes that can communicate directly. Even when all nodes are stationary, the network topology may be constantly changing due to variations in RF propagation and interference [6]. Thus, the network topology will likely be different when it rains than when it is sunny and different during a party than when a house is empty. A multi-hop network must adapt to the dynamically changing topology to allow nodes to communicate.

One approach to routing data in a multi-hop network is to have every node repeat every new packet received. This approach is advantageous in that it is simple, it routes between any two nodes, and it utilizes all redundant paths for greater reliability. However, because every node sends every packet once, this approach does not benefit from spatial re-use. Instead, this type of network “flooding” is typically used to identify a route through the network over which many data packets can then flow.

Network routes can be identified proactively or reactively. Proactive approaches maintain connectivity and resource availability information even when no traffic is present [20]. This approach reduces start-up latency, but wastes power and bandwidth when no routes are required and when the network changes frequently (in which case the information becomes stale). In contrast, a reactive routing approach identifies a route only after a packet transmission request or stream connection request is received [12] [21]. In networks with low utilization or highly dynamic topologies, a reactive approach is typically superior [5]. A hybrid

approach is possible in which some paths are maintained proactively while others are identified reactively. A network might also switch between proactive and reactive routing in response to network load.

Multi-Hop Channelization

Channelization is often used to alleviate the problems of single-hop and two-hop interference in a wireless network. The Request to Send (RTS)/Clear to Send (CTS) scheme, part of the Distributed Coordination Function (DCF) in WLAN standards, such as 802.11a or 802.11e [4], is one technique that may be used for this purpose. This technique allows nodes to acquire the channel and suppress other nodes from contending for the same channel. This, however, does not take advantage of multiple channels available in most standards [1] [2] [4]. These schemes allow nodes separated by two hops to communicate at the same time, if they choose non-conflicting channels. In a multi-hop network, such a channelization scheme may allow nodes to take advantage of multiple access features.




Figure 4: Colors representing channels (sub bands) are assigned in order to avoid interference

Solving this problem is akin to solving channelization based on the well-known graph coloring problem (Figure 4), which is known to be NP-complete¹. However, heuristic solutions may be implemented in the network using static, centralized, or distributed techniques. The small scale of home networks makes brute-force approaches possible.

QoS Routing and Multi-Hop Channelization

Routing with QoS may be implemented assuming that the underlying network supports a contention-free environment, derived through channelization schemes. However, if multi-hop channelization is completed a priori, without considering the needs of traffic, there is no way to guarantee the needed QoS (even using

¹ “No polynomial time algorithm has been discovered for this class of problem” [7].

RTS/CTS schemes). To maintain QoS in a dynamic network with varying application and usage scenarios, two problems must be addressed: resource management across the network and resource management at the node/link level.

1. *Resource management across the network—routing and channel assignment:* For a given connectivity graph of nodes and set of flows with a known link capacity requirement, channel resources (time and frequency) need to be allocated in an efficient manner. A choice made by one node will affect all other nodes in the network.
2. *Resource management at a single node and link level:* For a given set of flows entering and leaving a node and a known link capacity (assuming routing is complete and an end-to-end path is set up), the link must service flows to meet the QoS needs of each flow. At the outset, this problem seems similar to its wired counterpart. But new evidence suggests that wireless channel characteristics require special attention [16].

While separation of routing and channel assignment from QoS simplifies the path assignment problem, it is at a cost to the effective system capacity, as the flows are not known a priori. On the other hand, if we increase the complexity of the QoS path set-up problem, we can solve the channel assignment problem along with the QoS constraints (Figure 5). Even though the allocation in Figure 5 uses more sub bands than the one shown in Figure 4, the sub bands are allocated based on the QoS constraints. This approach results in tighter QoS guarantees and better overall network utilization.




Figure 5: Colors representing channels are assigned based on QoS requirements

Further performance advantages may be achieved by enabling nodes from neighboring networks to cooperatively provide QoS (Figure 2). For example, we might predict the capacity that would be available if only using nodes from a given network and allow QoS

reservations up to that capacity, but then allow messages to actually flow over any available nodes, thus taking advantage of extra performance. The remaining bandwidth would serve as leeway for the QoS algorithm or extra bandwidth for non-reserved uses (such as Web surfing).

These issues require the tradeoff of computation complexity to increase system utilization. This leads to the following questions:

1. What is the maximum size of the network to which the design of our algorithms should scale?
2. Should we jointly optimize the channel assignment and QoS path selection problems?
3. Should we assume all nodes have the same MAC and routing capabilities?
4. How do we allow new data traffic to preempt old data traffic?

Centralized Vs. Distributed Path Selection

Additionally, an important question in the choice of a QoS routing algorithm is whether to use a central controller or make decisions in a distributed manner. A decentralized QoS routing algorithm distributes the complexity across the network. This may increase the cost of nodes in the network, but distributing the decision making relaxes the need for network-wide synchronization and channel assignment, increasing the network scalability. A distributed approach also eliminates the need to communicate with a centralized controller, reducing the traffic overhead. Finally, while a decentralized approach can react more quickly to local changes, by eliminating the global-view of the network, it generally makes non-optimal decisions.

In a home networking environment, the relatively small network size and cost considerations are probably the most important factors to consider for making a decision between the two approaches. In such an environment, a centralized approach is more advantageous, as we can move most of the cost/complexity to the central nodes. The relatively small network size also implies that the overhead traffic flows are small, the reaction to the changes is quicker, and we may be able to make globally optimal decisions on the central controller.

Using a centralized method to solve problems of channelization, multi-hop routing, and path selection in a combined fashion will allow the system to meet the requirements of the Digital Home. Additionally, these algorithms need to comprehend security in order to protect the performance of the network. Not doing so will render any effort on their part useless in many of the deployment environments.

SECURITY

Security for home networking is a large topic covered more fully in “Home Network Security” also in this issue of the *Intel Technology Journal* [9]. This section mainly covers those security issues directly related to multi-hop networking. These issues were highlighted in the requirements section, namely, device introduction, user data/command integrity and secrecy, and network authentication.

Device and Router Introduction

In general, the introduction problem has been addressed by Universal Plug and Play (UPnP*) Security [9], and further work on simplifying introduction in the general case is being conducted as part of that R&D activity. This device introduction is a key component that enables solutions to the other security issues.

User Data and Command Security

User data security is not unique to multi-hop networks and is orthogonal to other security issues in these networks. The “Home Network Security” paper [9] in this issue of the *Intel Technology Journal* makes a case for end-to-end security. A summary of the motivation and its relationship to multi-hop networking is provided here.

Many homes require multiple security domains. For example, when the home contains older teenagers, adult roommates, boarders, or guests, we would expect each to have his or her own security domain, but these people would all share the same physical multi-hop network.

A security domain is a set of entities (devices or even processes) that are allowed to work together, excluding other entities. When traffic is encrypted for confidentiality, members of a domain can read one another’s traffic. When user-command traffic (to control DVD players, jukeboxes, etc.) is strongly authenticated and authorized, one member of a domain is permitted to issue commands to another. When a home has multiple security domains, security cannot be implemented at the link layer between two directly communicating devices; rather, it must be implemented as a protocol at a higher layer between the end points in a communication that may traverse multiple devices in a multi-hop network.

One advantage of a higher-layer end-to-end security protocol is that the physical network carrying the user’s traffic is not responsible for providing security for that traffic, and as a result, neighboring networks may be

used to help carry a user’s traffic without introducing security concerns.

Network Authentication

Even though user data security and integrity is provided end-to-end, devices still need to ensure network packets are received from authorized nodes in a multi-hop network. In particular, authentication must be supported for the one-hop origin of all packets, while the multi-hop originator must be authenticated for packets that access or control QoS or channelization on routers. A packet could be implicitly authenticated by encrypting it via a symmetric key known only to the origin and the verifier, but such a mechanism is more expensive than necessary and requires n^2 keys, where n is the number of network nodes. An alternative is to use a routing header in packets that can be authenticated with a lower computation overhead, resulting in a cost savings for routing nodes.

Routing and QoS Security

A routing algorithm makes decisions based on information such as latency and bandwidth between nodes. While the algorithm can acquire this information from neighboring network nodes, it is not possible to trust the routing information without verification. Without authentication, it is not possible to know if neighboring nodes belong to a possible adversary.

If a malicious node were used for routing, one way to interfere with the routing of messages is to insert an artificial delay in the message delivery path. It is not possible to prevent a malicious node from doing this. However, one can learn of this delay and route around it. This is therefore a normal routing problem rather than a routing security problem.

Another way to interfere with routing decisions is to advertise more bandwidth or lower latency for delivery of messages than can truly be achieved. One can test any advertised path for actual performance, provided that the destination node can authenticate a reply to a ping from the sending node. If we use public-key cryptography for this authentication, then we need a way to bind a node address to a public key. If the node address is the hash of the public key, then we achieve that binding without any additional cost.

Similarly, QoS establishment and maintenance requires nodes to trust information they learn from neighbors. QoS security cannot be free of administration. At the very least, a network node must learn which other nodes are part of its network. Clients will know their own QoS requirements and can make a reservation request from the nearest node belonging to the same owner. However, in a case of over-reservation, some human

* Other brands and names are the property of their respective owners.

administration will be required to establish priorities for different classes of use.

Solving the Coexistence Problem

Routing between neighboring networks (i.e., in adjacent apartments) may not be entirely independent and may require interference avoidance. Because security is provided end-to-end, the problem of coexistence between networks is only about sharing available bandwidth and not about maintaining privacy or data integrity. Three sharing strategies are possible:

1. Neighbors could choose to compete for channel access with no coordination between channel assignments. Such competition introduces channel contention. Channel contention reduces the overall channel capacity at high load and makes it difficult to predict the realizable channel capacity.
2. Neighbors could agree to statically split the available channels to avoid interference, thus introducing a constraint that must be reflected in QoS routing decisions. Channel assignment would be made independent of load, arbitrarily restricting the maximum bandwidth provided to each user wherever the networks overlap (typically along the common wall).
3. Finally, neighbors could agree to cooperate in channel assignment. While channel assignments could be initially made arbitrarily, one network could “borrow” a channel from the neighboring network (when not in use) or route traffic through a node on the neighboring network (i.e., nodes along a common wall) in response to high load. In this case, the network must be prepared to react if these resources are later reallocated by the neighboring network. One approach is to only allow neighboring resources to be borrowed for low-priority or non-QoS traffic. The cooperative approach decreases privacy, since neighboring networks must exchange load requirements to achieve QoS scheduling.

In each of these cases, the one-hop origin of packets containing routing updates or data from QoS-protected streams must be authenticated. We have described solutions for such authentication in the previous section.

RELATED WORK

The use of multi-hop wireless networking is gaining traction in everyday life. In fact, several companies already provide services using multi-hop wireless networks. MeshNetworks [17], for example, uses multi-hop routing between nodes installed on light poles, buildings, vehicles, and end-user devices such as laptops and handhelds to provide Internet access to subscribers

in cities. Nokia supplies kits to enable multi-hop networking between nodes installed on rooftops [18] to provide broadband Internet access. Most of these services focus on extending the reach of Internet access beyond the range typically supported by access points.

Multi-hop wireless networks exhibit many unique problems, but they also overlap with wired home networks. Device discovery and auto-configuration protocols such as Universal Plug and Play (UPnP*) [3] that were originally designed for wired IP networks can easily be applied to wireless networks. Wired home network security issues [9] also must be dealt with in wireless networks. Finally, QoS routing [19] and preemption have been extensively explored for wired networks.

QoS routing in ad hoc wireless networks has only recently been investigated in simulations. Several schemes were originally developed using MAC-independent techniques based on existing ad hoc routing protocols [14] [22]. More recently, advances have been made by optimizing QoS in multi-hop wireless networks with Code Division Multiple Access (CDMA) [15] and Time Division Multiple Access (TDMA) [24].

CONCLUSION

Multi-hop wireless technology offers unique benefits for creating a high-speed, robust home wireless network. The benefits over traditional infrastructure wireless networks include extending coverage without requiring deployment of multiple wired base stations, increasing utilization of spatial capacity to realize higher throughput, and offering alternate communication paths to provide failure recovery and better throughput.

To support the demanding usage models for the digital home, wireless networking innovations are required across the physical, MAC, and routing layers. In addition, higher level issues such as Quality of Service (QoS) guarantees, device discovery, and security must be solved. We have outlined a roadmap to meet these challenges; solving them will usher in new opportunities for wireless networking in the digital home.

ACKNOWLEDGMENTS

We acknowledge contributions from Stephen Wood, Sumit Roy, Rahul Mangharam and Arjunan Rajeswaran. We also thank reviewers of this paper for their valuable comments and suggestions.

* Other brands and names are the property of their respective owners.

APPENDIX A

Reconfigurable Radios

Support for a wide variety of consumer wireless PHY/MAC layers will be possible by using low-cost reconfigurable radios at the “edge” of the network. Furthermore, within the network it is highly advantageous, due to the varying Quality of Service (QoS) and bandwidth requirements, to have features like variable modulation [4] in the PHY to optimize the throughput based on channel loading and propagation conditions (party-time vs. home alone). So an example radio might have a very flexible PHY/MAC to interface externally, and it might use a common technique like 802.11, which has variable modulations, to route within the network itself.




Figure 6: Reconfigurable radio interfaces allow multi-hop and other wireless networks to bridge multiple PHY technologies

Reconfigurable radios should then be configurable to handle 802.11 a,b,g, low-rate consumer devices (lower requirement than 802.11) [25], and perhaps future high-data-rate (~100-500 Mbps) wireless personal area networks [1] and USB standards.

Radios should also be “intelligent,” assessing the environment and channel propagation around them, including interference. First, the radio should be able to broadcast and receive beacon signals introducing itself to the surrounding devices that are “reachable.” Next, the radio should be able to assess the existing frequency channels, interference, and noise conditions on each channel, so that it can determine (perhaps in conjunction with central or distributed control algorithms) the best channel to support a given bandwidth and QoS requirement. Having measured the channel, the radio can use the minimal power (to save power to the power amplifier) to make a reliable connection. Thus, unlike

conventional radios, a desirable radio requires a simple processor to help direct its various reconfigurable modes.

Finally, to lower cost, the reconfigurable radio (which by its very name implies higher cost) should take advantage of extensive silicon (Si) re-use. The baseband radio will have re-usable components such as various filters, digital mixers, etc., and the Radio Frequency (RF) portion will have a degree of agile frequency capability (for example, it will be able to jump to the 5.2 GHz band if there is microwave interference detected at the 2.4 GHz band). The MAC layer will also be flexible to allow download from the host of the typical extensive memory resources required to support the wide variety of anticipated MAC protocols: this will allow significant cost savings via memory re-use.

In summary, desirable home network radio architecture will require intelligence and reconfigurability to minimize power and keep QoS high while utilizing extensive Si re-use to keep costs down.

REFERENCES

1. “IEEE Wireless Personal Area Networks (WPAN) 802.15 High Rate (HR) Task Group (TG3).” <http://grouper.ieee.org/groups/802/15/pub/TG3.html>
2. “Specification of the Bluetooth System,” Ver. 1.1, February 22, 2001. <http://www.bluetooth.com>
3. “Universal Plug and Play (UPnP) Device Architecture Document,” Ver. 1.0, June 2000. <http://www.upnp.org>
4. “Wireless LAN Medium Access Control and Physical Layer Specifications,” Aug. 1999, IEEE 802.11 Standard (IEEE Computer Society LAN MAN Standards Committee).
5. J. Broch, D. A. Maltz, D. B. Johnson, Y. C. Hu, and J. Jetcheva, “A Performance Comparison of Multi Hop Wireless Ad-Hoc Network Routing Protocols,” in *Proc. of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM '98)*, Dallas, TX, October 1998.
6. D. Cheung and C. Prettie, “A Path Loss Comparison Between the 5 GHz UNII Band (802.11a) and the 2.4 GHz ISM Band (802.11b),” Intel Corp. Technical Report, January 2002. http://impulse.usc.edu/resources/802_11a-vs-b_report.pdf
7. T. H. Cormen, C. E. Leiserson, R. L Rivest, *Introduction to Algorithms*, McGraw-Hill, 1998.
8. S. Diaz, “Wireless networking: daunting but doable,”

- San Jose Mercury News*, May 29, 2002.
<http://www.siliconvalley.com/mld/siliconvalley/3363171.htm>
9. C. Ellison, "Home Network Security," *Intel Technology Journal* Vol. 6, Issue 4, 2002.
 10. J. Geier, "802.11a/b Site Survey: A Testimonial," *802.11 Planet*, October 10, 2002.
<http://www.80211-planet.com/columns/article.php/1479831>.
 11. P. Gelsinger, Intel Developer Forum Keynote, Feb. 28, 2002.
<http://www.intel.com/pressroom/archive/speeches/gelsinger20020228.htm>
 12. D. B. Johnson, "Dynamic Source Routing in Ad Hoc Wireless Networks," in *Mobile Computing* (ed. T. Imielinski and H. Korth), 1996.
 13. J. Kosseff and E. Hand, "Wireless channel use sets up turf battle," *The Oregonian*, August 19, 2002.
 14. S. Lee and A. T. Campbell, "INSIGNIA: In-band signaling support for QoS in mobile ad hoc networks," in *Proc. of the 5th International Workshop on Mobile Multimedia Communications*, 1998.
 15. C. R. Lin, "On-demand QoS routing in multihop mobile networks," in *Proc. of INFOCOM 2001*, Anchorage, AK, April 2001.
 16. R. Mangharam, "HotSpot: Access Point Services," Technical Presentation, July 24, 2002.
 17. Mesh Networks, Inc.
<http://www.meshnetworks.com/>
 18. Nokia Rooftop Solution.
<http://www.wbs.nokia.com/solution/>
 19. C. Parris, H. Zhang, D. Ferrari, "Dynamic Management of Guaranteed Performance Multimedia Connections," *ACM Springer-Verlag Multimedia Systems Journal*, Vol. 1, No. 6, 1994.
 20. C. E. Perkins and P. Bhagwat, "Highly dynamic destination sequenced distance vector routing (DSDV) for mobile computers," in *Proc. of the SIGCOMM'94 Conference on Communication Architectures, Protocols and Applications*, August 1994.
 21. C. E. Perkins, E. M. Royer, and S. R. Das, "Ad-hoc On-Demand Distance Vector (AODV) Routing," IETF draft. draft-ietf-manet-aodv-06.txt.
 22. E. M. Royer, C. Perkins, S. R. Das, "Quality of Service for Ad Hoc On-Demand Distance Vector Routing," IETF Draft, draft-ietf-manet-aodvqos-00.txt, Work in Progress, July 2000.
 23. L. A. Rusch, "Indoor Wireless Communications: Capacity and Coexistence on the Unlicensed Bands," *Intel Technology Journal*, Q3 2001,
 24. C. Zhu and M. S. Corson, "QoS routing for mobile ad hoc networks," in *Proc. of the Twenty First International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, New York, NY, June 23-27, 2002.
 25. ZigBee Alliance. <http://www.zigbee.com/>

AUTHORS' BIOGRAPHIES

Lakshman Krishnamurthy is a Senior Staff Engineer in the Network Architecture Lab, Intel R&D. His current research interests include high-speed wireless networking, wireless sensor networks, and software-defined radios. He received B.E. and Ph.D. degrees from the University of Mysore, India and the University of Kentucky, respectively. His e-mail is Lakshman.Krishnamurthy@intel.com.

Steven Conner is a Senior Network Software Engineer in the Network Architecture Lab, Intel R&D. His current research interests include wireless networking and network self-configuration. He received B.S. and M.S. degrees from the University of Arizona. His e-mail is w.steven.conner@intel.com.

Mark Yarvis is a Senior Researcher in the Network Architecture Lab, Intel R&D. His current research interests include techniques for leveraging network heterogeneity, sensor networks, and pervasive computing. He received B.S., M.S., and Ph.D. degrees from the University of California, Los Angeles. His e-mail is Mark.D.Yarvis@intel.com.

Jasmeet Chhabra is a Senior Network Software Engineer in the Network Architecture Lab, Intel R&D. His current research interests include wireless sensor networks and software-defined radios. He received a B.E. and M.S. degree from the University of Delhi, India and the University of Maryland, College Park, respectively. His e-mail is Jasmeet.Chhabra@intel.com.

Carl M. Ellison is a Senior Security Architect in the Network Architecture Lab of the Corporate Technology Group of Intel Corporation. His current research is devoted to delegatable, distributed, public-key authorization with a special emphasis on self-organizing networks. His concentration on security has been a side effect of a more general career focus on distributed and fault tolerant systems. His e-mail is cme@jf.intel.com.

Chuck Brabenac is a senior architect in the Communications and Interconnect Technology Lab, Intel R&D. He is currently involved in ultra wideband (UWB) wireless communications research, with focus on its associated MAC and application stacks. He is also active in UWB-related standards activity in IEEE 802.15.3a, where he will be serving as vice-chair of this task group that is chartered to develop a high data rate WPAN standard. His e-mail is chuck.brabenac@intel.com.

Ernest Tsui is a principal engineer and manager of the Wireless Architecture Research group, Communications and Interconnect Technology Lab, Intel R&D. His interests at Intel are in wireless (and wireline) communications algorithms and architectures. He is owner of the reconfigurable baseband PHY architecture for CTG. He received his B.S., M.S., and Ph.D. degrees from U.C. Berkeley and is a Senior Member of the IEEE. His e-mail is ernest.tsui@intel.com.

Copyright © Intel Corporation 2002. This publication was downloaded from <http://developer.intel.com/>.

Legal notices at <http://www.intel.com/sites/corporate/tradmarx.htm>.

For further information visit:

developer.intel.com/technology/itj/index.htm