An Inter-domain Routing Protocol for Multi-homed Wireless Mesh Networks

Yair Amir[1], Claudiu Danilov[2], Raluca Musăloiu-Elefteri[1], Nilo Rivera[1]
[1] Johns Hopkins University
{yairamir, ralucam, nrivera}@dsn.jhu.edu
[2] Boeing Phantom Works
claudiu.b.danilov@boeing.com

Abstract

This paper presents an architecture and a hybrid routing protocol for multi-homed wireless mesh networks that provide uninterrupted connectivity and fast handoff. Our approach integrates wireless and wired connectivity, using multicast groups to coordinate decisions and seamlessly transfer connections between several Internet gateways as mobile clients move between access points. The protocol optimizes the use of the wireless medium by short-cutting wireless hops through wired connections, paying a very low overhead during handoffs. The paper demonstrates that inter-domain handoffs occur instantaneously, with virtually no loss or delay, for both TCP and UDP connections.

1. Introduction

Wireless mesh networks extend the connectivity range of mobile devices by using multiple access points, some of them connected to the Internet, to create a mesh topology and forward packets over multiple wireless hops. Mobile clients can freely roam within the area covered by the mesh access points and maintain their connectivity at all times.

As the size of a wireless mesh network increases, the number of Internet connected access points (*Internet gateways*) needs to increase to disperse traffic and avoid congestion. In practice, Internet gateways will reside at different locations and will often be connected to different network domains. We refer to such mesh networks as *multi-homed*. In this type of networks, a mobile client is served by a nearby access point that forwards data packets (potentially over multiple wireless hops) to its closest Internet gateway.

Multi-homing poses a challenge in providing continuous connectivity to mobile clients that may move between the areas covered by different access points. Those access points will often have different Internet gateways closest to them. When such a transition (*handoff*) occurs, we would like to maintain all previously opened connections, and transfer them to the new Internet gateway as quickly as possible, without any involvement from the mobile device.

This paper presents a simple and elegant architecture that supports seamless routing in multi-homed wireless mesh networks. The routing protocol integrates wired and wireless communication and optimizes performance of the hybrid routing, in our case by minimizing the usage of wireless transmissions. The handoff between Internet gateways is completely transparent to the mobile devices, which have no indication of when, or whether a handoff takes place at all, and is fast enough for real-time applications, such as VoIP, where any interruption in connectivity can have an adverse impact on the service quality. While solutions exist for intra-domain handoff [5, 11, 15], we believe that currently there are no efficient and transparent protocols for inter-domain handoff in multi-homed wireless mesh networks.

In our approach, new connections always use the closest Internet gateway at the time of their creation, while existing connections are forwarded through the wired infrastructure to the Internet gateway where they were originally initiated. As the handoff process requires routing agreement and transferring connections between the involved Internet gateways, our protocol guarantees that packets are routed correctly, at all times.

We implemented our protocol on the SMesh [5] system in order to support optimized hybrid wireless-wired routing and fast inter-domain handoff. SMesh is a seamless wireless mesh network that provides intra-domain handoff with real-time performance. We believe that the combination of the inter-domain routing protocol presented in this paper, and the intra-domain connectivity provided by SMesh, realizes the first complete multi-homed wireless mesh network that is transparent to mobile devices, and provides a fast intra- and inter-domain handoff suitable for real-time applications such as VoIP. The system is currently deployed over three buildings at the Johns Hopkins University campus, is used by several students and faculty on a daily basis, and

1-4244-0992-6/07/$25.00 © 2007 IEEE

the software is freely available to the community.

The main contributions of this paper are:

- A simple and practical architecture that integrates seamlessly wired and wireless connectivity in multi-homed wireless mesh networks.
- A hybrid routing protocol for wireless mesh communication that optimizes routes as mobile devices move between Internet gateways.
- A fast inter-domain handoff protocol for multi-homed wireless mesh networks that supports real-time applications such as VoIP.

The rest of the paper is organized as follows: Section 2 presents related work. In Section 3 we describe the architecture of our multi-homed mesh network approach, and in Section 4 we present the real-time handoff protocol between the wired Internet gateways. Experimental results are presented in Section 5, and Section 6 concludes the paper.

2. Related Work

We propose a mechanism for multi-homed wireless mesh networks that optimizes routing and provides fast inter-domain handoff between Internet connected access points, potentially on different networks. As such, our work relates to mobility, wireless mesh networks and wireless handoff. Good surveys addressing some of these areas were overviewed by Akyildiz et al. in [4] and [3]. Note that related work may also refer to intra-domain handoff as *micromobility* and to inter-domain handoff as *macromobility*.

Two general approaches that can support intra-domain and inter-domain handoff are Mobile IP (MIP) [13] and Mobile NAT [7]. In MIP, a client binds to an IP address at the Home Agent (HA). As the mobile client moves to a different access point or domain, it receives a Care-of-Address (CoA) from a Foreign Agent (FA). The mobile client then registers its new CoA with its HA, and data is then tunneled through the HA. Some enhancements have been proposed to improve intra-domain handoff latency [9, 10]. Our approach does not require binding the mobile client to a specific Home Agent, but rather ties each connection to the Internet gateway that is closest at the time the connection is initiated.

In Mobile NAT, a client receives two IP addresses through DHCP: a binding address for the network stack, and a routing address that will be visible in the network. As the mobile client moves to a different domain, the client may receive a new routing address. However, as end-to-end connections were initiated from the IP address of the network stack, which remains the same, existing connections will be maintained. The approach requires modifying the mobile client network stack to be aware of the protocol, and also changes in the standard DHCP protocol. Our approach does not require any modifications to the mobile client thus supporting standard mobile devices of any architecture or operating system.

Existing experimental wireless mesh testbeds that support client mobility include MeshCluster [11] and iMesh [15], both of which work with mobile clients in infrastructure mode. MeshCluster, which uses MIP for intra-domain handoff, shows a latency of about 700 ms due to the delay incurred during access point re-association and MIP registration. iMesh also offers intra-domain handoff using regular route updates or MIP. Using layer-2 handoff triggers (no moving client), handoff latency in iMesh takes 50-100 ms. SMesh [5] provides 802.11 link-layer and network-layer fast handoff by working in ad-hoc (IBSS) mode, controlling handoff from the mesh infrastructure, and using multicast to send data through multiple paths to the mobile client to deal with incomplete knowledge and unpredictable moving patterns.

None of the above systems optimize routes on multi-home wireless mesh network. Also, none of them provide fast inter-domain handoff. In this paper, we use SMesh for providing mobile client transparency and real-time intra-domain handoff, and propose novel mechanisms that optimize wireless communication through the wired connections and provides a practical real-time inter-domain handoff between the mesh Internet gateways, thus providing the first complete solution for transparent multi-homed mesh networks with fast handoff.

3. A Hybrid Overlay Architecture

A wireless mesh network is comprised of multiple access points, possibly distributed in several islands of wireless connectivity such as different buildings located close to each other or parts of the same building. Access points inside a wireless island can communicate, potentially using multiple intermediate hops. One or more access points in each wireless island is connected to the Internet through a wired network. We call these access points *Internet gateways*. For Internet connectivity, other access points rely on multi-hop communication to reach an Internet Gateway in their island. Figure 1 shows an example of a wired-wireless hybrid mesh network with two islands, each of them with two Internet gateways.

Each access point runs a software router that allows multi-hop communication. These routers create an overlay topology where some of the links are wireless (between nodes in the same island) while others are wired (between the Internet gateways). There are several options available for a practical software router, including X-Bone [14], Spines [2] and RON [6]. In our implementation we use

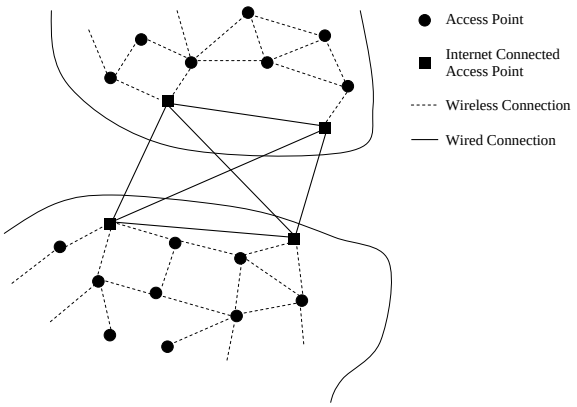

Figure 1: Hybrid Overlay Mesh Network

the Spines overlay messaging system to provide multi-hop communication as it offers overlay multicast, anycast and unicast forwarding. We make use of overlay multicast to auto-discover Internet gateways and to coordinate decisions between access points during mobile client handoffs. We use anycast to forward data packets from a client to the closest Internet gateway.

Using one overlay network for both wireless and wired communication has several advantages. Peer-to-peer communication between access points located in the same wireless island can take advantage of wired connectivity between remote Internet gateways to shortcut multiple wireless hops. In addition, the diameter of the network is decreased, improving route update latency and overhead related to control messages on overlay network.

3.1. Topology Formation

The topology formation starts with each access point broadcasting its presence periodically. Neighboring nodes create bidirectional links and advertise their connectivity through a link state protocol to other nodes in the network. The link state protocol uses link-based acknowledgments such that after a link was advertised to other access points in the network, it will not be advertised again, unless it changes its status. This reduces communication overhead for managing the topology.

Internet gateways join a multicast group called *Internet Gateway Multicast Group* (IGMG) on which they periodically advertise their wired interface IP address. The multicast routing is handled by the underlying overlay infrastructure. Multicast trees are calculated in a way similar to that of MOSPF [12]. When two Internet gateways receive each other's advertisements (which initially travels through the wireless infrastructure to the members of the multicast group), they connect through a wired overlay link. This way, the Internet gateways inside an island form a fully connected graph using their wired infrastructure, while the other access points inside the island interconnect based on the wireless connectivity. In order to interconnect wireless islands, at least one Internet gateway in each island needs to be pre-configured to connect to a set of Internet gateways such that an initial connected graph is formed. Then, multicast advertisements from all gateways will be propagated, Internet gateways will connect to each other, and eventually, a fully connected logical graph between all Internet gateways in all islands is formed.

3.2. Routing Metric

In a multi-homed wireless mesh network, some of the access points have wired connections that can be used to shortcut several hops of wireless communication, thus decreasing the number of wireless transmissions. In general, in a combined wired-wireless routing metric scheme, it is reasonable to assume that a wired connection costs much less than a wireless link. On the other hand, depending on the network conditions it is possible that wired connections between Internet gateways have different costs (based on throughput, loss rate, latency, etc.).

Our approach uses the best route to a destination considering wireless connectivity as well as any hybrid route available, and allows for different routing metrics to be used both on the wired and wireless links. Considering that each wireless link can have an $ActualCost$ metric of at least 1, the routing cost of that link will be: $Cost = ActualCost * (M + 1)$ where M is the maximum cost that can be associated with a wired path. For example, if a wired link can have a maximum cost of 10, and there are 5 access points connected to the Internet in the mesh network, the value of M is 40 (the largest number of wired hops in a path is 4), and the minimum cost of a wireless link is 41. The cost of a hybrid path is the sum of the cost of all the links. This mechanism gives preference to any wired link over a wireless one, and optimizes the wired path based on a desired metric. For example, we can use ETX [8] as the wireless $ActualCost$ metric, and latency as the wired links metric.

3.3. Handling Mobile Clients

Mobile clients connect to their closest access point and use it transparently as they would work with a regular Internet connected access point. No special software or drivers need to be installed on the mobile clients. The mesh network is responsible to forward packets to and from other clients or the Internet. In our implementation, all access points use a private IP domain (10.x.y.z) for their wireless interfaces. Mobile clients are assigned IP addresses through DHCP from the same IP domain.

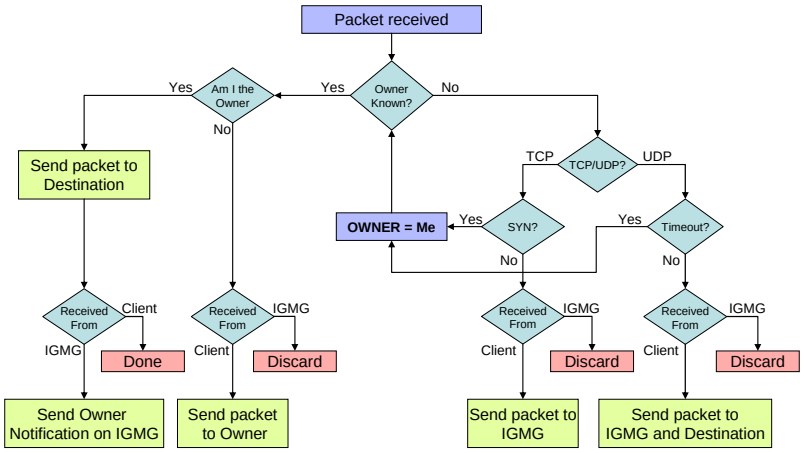

Figure 2: Inter-domain Handoff Flowchart

Mobile clients advertise their presence periodically by broadcasting DHCP requests to renew their leases. Access points that hear these requests start monitoring the client and use the overlay infrastructure to join a multicast group specific to the client's IP address. In our implementation, for a client with IP address 10.x.y.z, the access points in its vicinity join the group 224.x.y.z, called *Control Group*. Using this multicast group access points locally advertise their link quality. When an access point considers serving a mobile client, it joins a different multicast group, called *Data Group*, in our case 225.x.y.z. At this point, the access point sends a gratuitous ARP to the mobile client, associating its default route (next hop) with the MAC address of the new access point.

Packets sent to a mobile client are routed by the messaging infrastructure to the Data Group corresponding to the receiver client. Local access points that joined the Data Group then forward the packets to the mobile client. The reason for using a multicast group instead of a single IP address for the client packets is that in periods of instability, when it is not yet decided which local access point should serve the client, multiple access points in the vicinity of the mobile client may forward the data packets (also allowing us to deal with unpredictable moving patterns). When an access point receives a packet that has a destination outside the wireless mesh network, it simply forwards it to the Internet Gateway Anycast Group, an overlay anycast group to which all Internet gateways join. This way, packets are always sent to the closest Internet gateway. More details about intra-domain handoff and routing in SMesh can be found in [5].

4. Inter-domain Handoff Protocol

Packets exchanged between two mobile clients, either in the same or in different wireless islands, simply use shortest path multicast trees reaching the access points that decided to serve each client. Note that in the stable case, when mobile client communication does not require a handoff, only one access point in the vicinity of a client will join its multicast Data Group. Therefore, most of the time, the multicast trees are simply linear paths. The multicast trees adjust automatically when mobile clients roam within the vicinity of different access points, as the access points join or leave the client's multicast Data Group. In peer-to-peer communication, packets will follow the shortest paths with no need for a special handoff at the Internet gateways.

In contrast, communication between mobile clients and the Internet is relayed through the closest Internet gateway. As mobile clients move within the wireless mesh network, they may get closer, network-wise, to a different Internet gateway in the same island, or they may move to a different wireless island. In this case, the anycast packets, which are forwarded to the closest Internet gateway, will no longer reach the original gateway, and therefore a solution is required to maintain existing connections.

Mobile clients in SMesh work on a private network, and a Network Address Translation (NAT) is required at the Internet gateway when communicating with an external host. Each Internet gateway has a different external IP address. Applications using TCP, and in some cases, applications running on top of UDP require packets to be forwarded through the initial forwarding Internet gateway through the entire life of the connection. Changing one end-point of the connection (the IP address of the Internet gateway) is often impossible without breaking the existing connection, and therefore it is better for the handoff mechanisms to mask

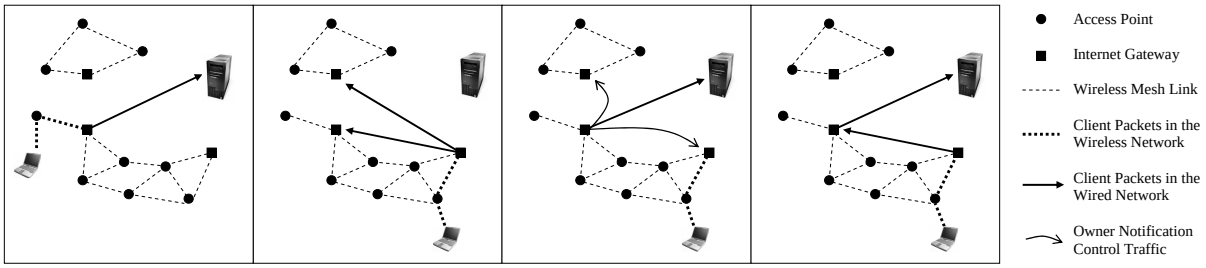

Figure 3: TCP forward handoff: (a) Connection establishment (b) Handoff Phase 1 (c) Handoff Phase 2 (d) Handoff completed

this problem inside the mesh network.

One potential solution is to exchange complete connection information (NAT tables) between the Internet gateways periodically and forward packets to the original owner of the connection using the wired connectivity. Such a solution can only be as fast as the time between two periodic NAT table exchanges, and cannot support real-time traffic such as VoIP. To support real-time traffic, one can advertise connection information to all the Internet gateways when the NAT entries are created. However, this technique tends to be wasteful, as not all mobile clients may move and change their Internet gateway. The problem is most notable when clients are browsing the Internet, as many connections are established for each website and, all of these information, which is relevant only for a small amount of time, would be sent to all of the Internet gateways.

Our inter-domain handoff protocol provides transparent mobility on a NATed network with real-time performance. We treat UDP and TCP connections separately, detect the existing owner (the Internet gateway from which the connection was initiated) of a connection, and forward existing connections through their original owners[1]. Figure 2 shows the general flow of packets at each Internet gateway.

4.1. TCP Connection Handoff

A TCP session requires that the source and destination IP addresses and ports to remain constant during the life of the connection. Our mobile clients run in a NAT address space, and although connections are end-to-end, the Internet destination regards the source address as that of the Internet gateway that sent the first SYN packet. When a mobile client moves closer to a different Internet gateway, the new gateway must forward all packets of each existing connection to the original gateway that initiated that connection. On the other hand, new connections should use the Internet gateway that is closer to the client at the current time, and not be forwarded to an old gateway.

In TCP, a SYN packet indicates the creation of a connection and generates a NAT entry, while a FIN packet indicates the destruction of the connection. If an Internet gateway receives a TCP packet that is not a SYN and it does not have an entry for that connection in its NAT table, it forwards that packet to the IGMG group. The original owner of the connection (the one that has it in its NAT table) relays the packet to the destination, and sends a message to the IGMG group, indicating that it is the connection owner for that NAT entry. Then, any gateway that is not the connection owner, will forward packets of that connection to the respective owner, finalizing the connection handoff process. Figure 3 shows the stages of such a TCP connection handoff.

If packets arrive at an Internet gateway at a fast rate, several packets may be sent to the IGMG group before the connection owner can respond. If no Internet gateway claims the connection within a certain timeout (in our implementation 3 seconds), the new gateway claims the connection, forwarding the packets directly to the Internet destination. This will break the TCP connection, which is the desired behavior in such a case, since it is likely that the original owner crashed or got disconnected. Causing the Internet host to close the connection avoids connection hanging for a long period of time (TCP default is 2 hours).

4.2. UDP Connection Handoff

Most real-time applications use the best effort UDP service and build their own protocol on top of UDP to meet specific packet latency requirements. Some applications, such as DNS, do not establish connections between participants. Others, such as SIP in VoIP, establish specific connections defined by a pair of an IP address and a port at both ends of the connection.

When an Internet gateway receives a UDP packet with a new pair of source and destination addresses or ports, it cannot distinguish between the case where this is the first packet of a new connection, and the case where the packet

[1]One can potentially spoof the address of the original owner to reduce the routing overhead of our protocol. However, egress filtering is commonly used at network routers and will prevent spoofed packets from leaving their network.

belongs to an existing connection established through a different Internet gateway.

We classify UDP traffic on a port number basis as *connection-less* and *connection-oriented*, and choose connection-oriented as the default protocol. Connectionless UDP traffic is forwarded directly after receiving it from the mesh network, on the current shortest path. DNS and NTP traffic falls into this category.

Upon receiving a new connection-oriented UDP packet that has an Internet destination, an Internet gateway relays that packet to its destination, and also forwards it to the multicast group that all Internet gateways join (as opposed to the TCP case, where the access point only sends packets to the multicast group). If the UDP packet belongs to a connection that was already established, the Internet gateway that is the original owner of the connection also relays the packet to the destination, and sends a response to the Internet gateway multicast group. After receiving the response, the initial gateway will forward subsequent packets directly to the original gateway, and will no longer relay UDP packets of that connection (with the same source and destination addresses and ports) to the Internet. If a response does not arrive within a certain timeout (in our implementation 500 milliseconds), the Internet gateway will claim ownership of the UDP connection, will stop forwarding packets of that connection to the IGMG group, and will continue to relay packets to the Internet.

4.3. Discussion

Due to handoff and/or metric fluctuations, there is a possibility that packets coming from a mobile client and belonging to the same flow alternate between two Internet gateways. This may lead to more than one gateways claiming the ownership of the connection. We encounter such case in TCP when a client retransmits a SYN connection request, and this request is routed through a different Internet gateway. In UDP, such case may occur when two different Internet gateways start forwarding client packets for the same connection at about the same time. A plausible solution for TCP is to delay ownership decision until a full three-way TCP handshake is seen by the Internet gateway. For UDP, when there is more than one ownership request in parallel, the gateways decide the rightful owner of the connection based on feedback traffic from the end-host or lowest IP address.

Also note that, in general, our inter-domain handoff protocol can be applied in less sophisticated architectures. For example, all Internet gateways can be pre-configured with the complete set of Internet gateways that will participate in the inter-domain handoff. However, route optimizations provided by the overlay network, both in the wired and wireless network, will not be available, and some other mechanism must be devised to ensure fast seamless handoff for mobile clients at the intra-domain level.

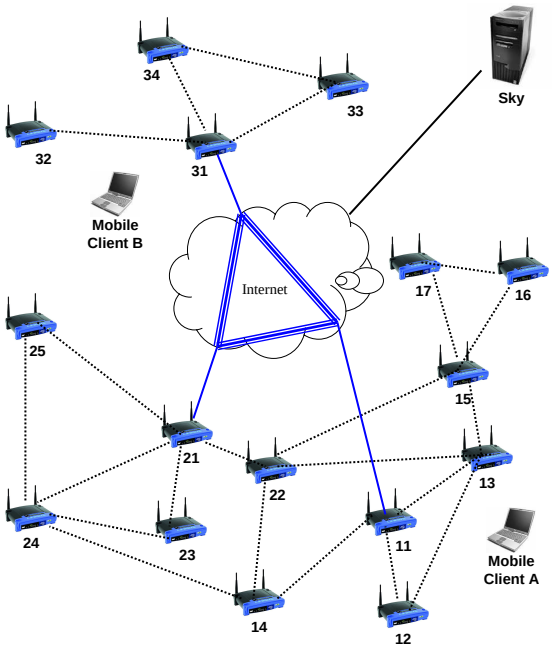

Figure 4: The Hybrid Wired-Wireless Testbed.

5. Experimental results

We implemented our protocols within the firmware of Linksys WRT54G wireless routers. A third party firmware (OpenWRT [1]) was installed in the Linksys routers to provide us with a Linux environment suitable for running our prototype. Other than adding our system executables, no other changes were made to the firmware.

We deployed our system on 16 Linksys WRT54G wireless routers across several floors in three buildings. Each of the routers is equipped with one radio configured in ad-hoc mode. Three of the routers were connected to the Internet. Transmit power of the access points was set to $50mW$. We used two Windows XP laptop computers with a Broadcom 802.11g Mini-PCI card in ad-hoc mode as the mobile clients. No software other than the benchmarking programs was installed on the laptop computers.

The topology of the wireless testbed used in our experiments is shown in Figure 4. The topology consists of one main island with two Internet gateways, and another smaller island with one gateway. The islands are disconnected due to a large open grass area between the buildings. However, a mobile client located between the two islands can reach both networks. Each of the Internet gateways is part of a different domain on the campus network and within 6 hops of each other through the wired network. Unless otherwise

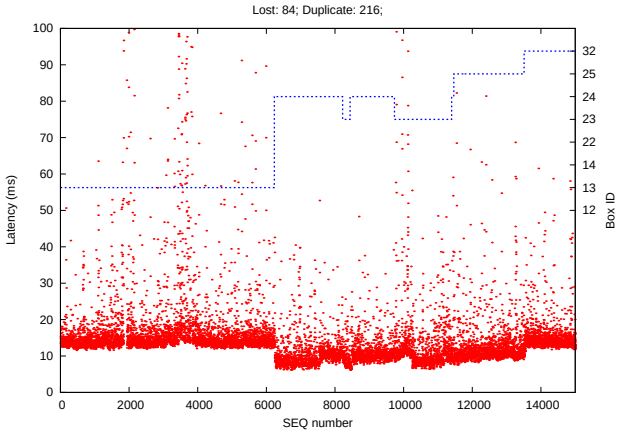

Figure 5: P2P test. Latency of packets received at Client A.

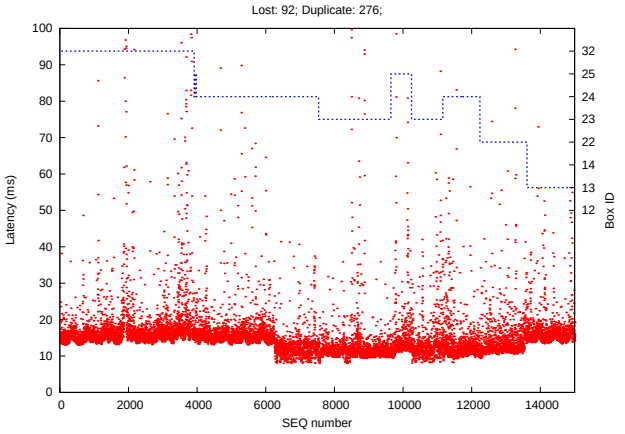

Figure 6: P2P Test. Latency of packets received at Client B.

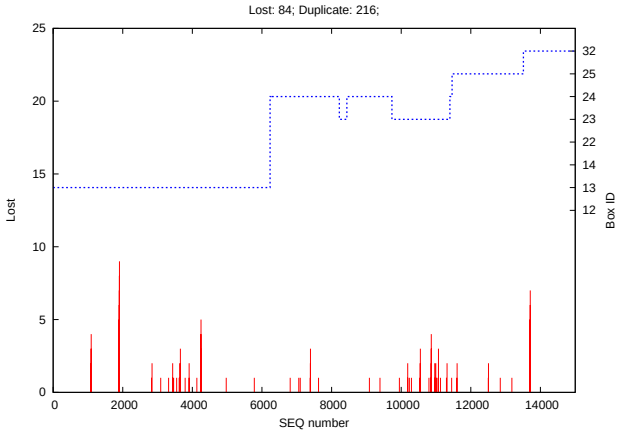

Figure 7: P2P Test. Lost packets at Client A.

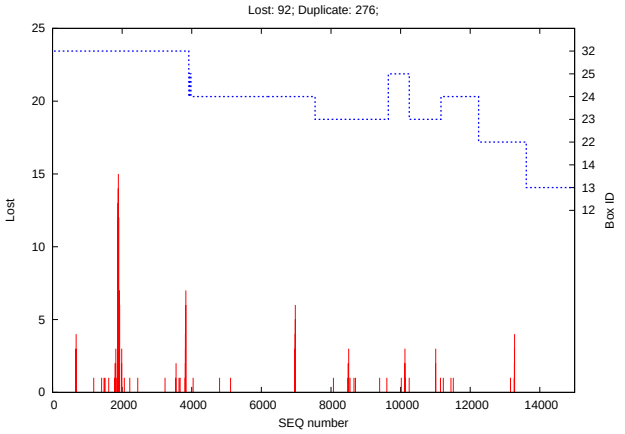

Figure 8: P2P test. Lost packets at Client B.

specified, the topology between the access points was static during the experiments. Each access point box has an identifier (box-id). The box-id of Internet gateways ends with digit 1 (Boxes 11, 21, and 31). The closest Internet gateway of mesh nodes is given by the prefix of the access point box-id (i.e. Box 23 uses Box 21 as its Internet gateway).

Experiments consist of walking with a mobile client from the 3^{rd} floor of a building located in the main island to a hallway in the second floor, followed by going down to the ground floor. Then, while walking outside on an open grass area we end up reaching the second island. This movement results in a few access point handoffs and at least three Internet gateway handoffs. A mobile client will be referred to as *Client* and the Linux box from the Internet as *Sky*. In all experiments we send a full-duplex (two-way) VoIP traffic. The VoIP traffic consisted of 160 byte packets sent every 20 ms at a rate of 64 Kbps, for 5 minutes. We focus our experiments on VoIP as a representative application that poses severe latency requirements.

Peer-to-peer UDP test: During this experiment two mobile clients walk in opposite directions from different buildings towards the original position of the other mobile client. Routing decisions are based on the path that decreases the number of wireless hops between the clients in the hybrid wired-wireless overlay network. Figures 5 - 8 present the results of this experiment.

In each graph, the access point that serves the mobile client is shown on the right vertical axis. The current access point is represented with a continuous dotted line. Horizontal plateaus of the dotted line represent stable periods in which the access point serves the client, while vertical jumps between plateaus represent handoffs between access points. For example, Figure 5 shows a transition from Box 13 to Box 24 around packet number 6000.

In this experiment, *Client A* started from the island which forwards traffic through the Internet gateways 11 and 21 while *Client B* started from the other island, which forwards traffic through gateway 31. Figures 5 and 6 show the one-way latency of packets as they are received at each client. The initial latency represents 4 wireless hops plus 1 wired hop. This is because there is one wireless hop between each client and its access point, and both access points are one wireless hop away from their corresponding gateway.

First note that, at around packet 4000, Client B handoff

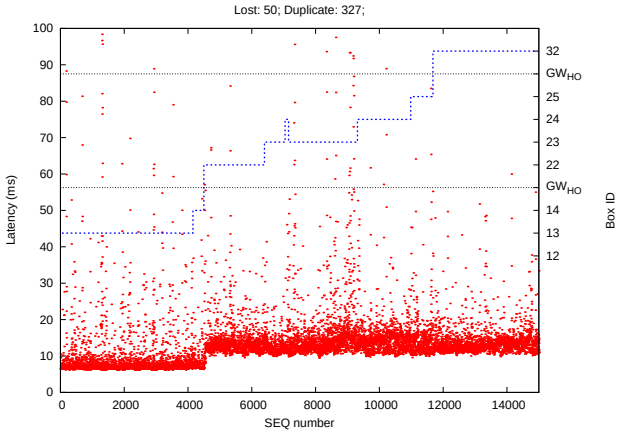

Figure 9: Inter-domain test. Client is receiver. Latency.

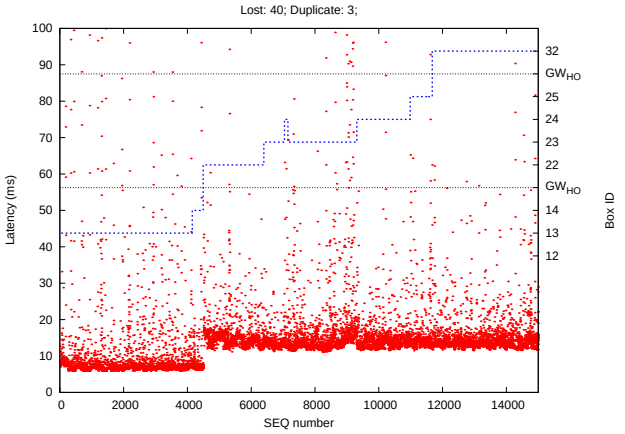

Figure 10: Inter-domain test. Sky is receiver. Latency.

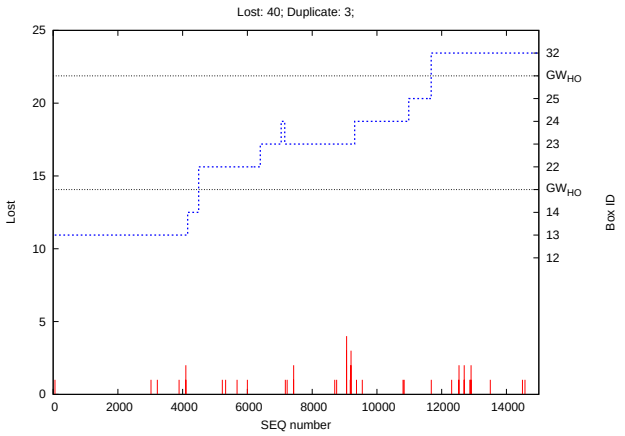

Figure 11: Inter-domain test. Sky is receiver. Loss.

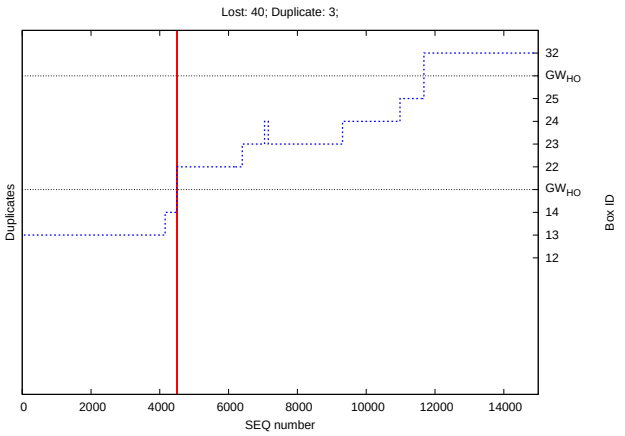

Figure 12: Inter-domain test. Sky is the receiver. Duplicates.

to access point 24, which makes Clients A and B reside in the same island. After this handoff, Clients A and B continue to use the wired network through the closest gateway to minimize the wireless usage. That is, if the clients were to only send through the wireless network, they would use 5 wireless hops instead of 4 and the wire. The latency decreases shortly after packet 6000 when Client A handoff to Box 24. At that point, both clients send and receive data through the same access point. Latency continues to change depending on the client's access point and the number of wireless hops between them. Just before packet 14000, the latency goes back to the latency at the beginning of the test since both clients are again four wireless hops plus one wired hop away from each other.

Overall, 84 packets were lost in one direction and 92 in the other. Figures 7 and 8 present the packets lost at the two clients. Loss is represented as cumulative number of losses over the last 25 packets. For example, Figure 8 shows the highest cumulative loss around packet 2000, where 15 consecutive packets were lost. As the wireless medium is shared, a sudden loss may be triggered by a number of factors including external wireless communication or interference from our own wireless network. In most real time applications, the effect of a relatively small number of packets being lost can be compensated with no interruption in service or significant quality degradation.

Connection Oriented UDP test: This test is done between a single mobile laptop, *Client*, and the Internet connected machine, *Sky*. Figures 9 and 10 show the one-way packet latency for packets received at *Client* and *Sky*, respectively. The horizontal lines marked `GW_HO` separate the graph into three areas defined by the Internet gateway forwarding the mobile client's packets to and from the Internet. An inter-domain handoff happens when the dotted line, showing the current access point serving the client, crosses one of the horizontal line.

Both latency graphs show a jump of around 4 ms after the first Internet gateway handoff. This is due to forwarding packets between gateways through the wired overlay network. We also note that occasionally, some of the packets have a higher latency just before a handoff (e.g., around sequence number 9000). This is due to 802.11 retransmissions, and is expected to happen when a mobile client moves away from its current access point. After the hand-

off, a better connected access point starts serving the client and the packet latency decreases.

Figure 11 shows the packets lost at *Sky*. We can see that losses occurred in bursts of less than 5 packets. The number of packets that arrived after more than 100 ms was 18 in the stream from *Sky* to *Client* and 92 in the stream from *Client* to *Sky* (which are considered lost in VoIP). Considering the total number of packets (15000 in each direction), very few packets were lost or delayed.

In Figure 12 we show the duplicate packets received by *Sky*. These duplicates are caused by inter-domain handoffs. There were only 3 duplicate packets on the stream in the entire experiment, and they occurred during the first Internet gateway handoff. Since Box 21 was not aware initially whether the packets belong to a new or an already existing connection, it sent the traffic both to the IGMG group and to the final destination (as explained in Section 4.2). Because Box 11 already had a connection established for that stream in its NAT entries, it forwarded the packets to the Internet destination, and at the same time, it notified the other gateways that it is the owner of the connection, by sending an acknowledgment to the IGMG group. As soon as Box 21 received an ownership acknowledgment from Box 11, it stopped relaying packets to *Sky* and start forwarding the packets to Box 11. Note that after the notification, all gateways learned about the ownership of that connection. This is the reason there are no duplicates in the second gateway handoff, from Box 21 to Box 31 that occurs before packet 12000.

TCP test: Our TCP tests show similar results to those presented on the previous UDP test. Also, connections did not break when moving around the mesh. However, the latency of packets tended to be higher for some packets. For example, the number of packets that arrived after 100 ms was 2 to 3 times higher for similar experiments. The main reason is that TCP delays packets that arrive out of order (mainly due to lost packets in the wireless) until it recovers the losses and can deliver the packets in order.

Mesh Gateway Failure test: It is interesting to see what happens when the Internet gateway used by a TCP connection suddenly fails. If that Internet gateway is the owner of the connection, then we expect that the connection will break. However, if the Internet gateway is not the original owner of the connection, but rather the one closer to the mobile client that forwards packets to the owner Internet gateway, we expect the mesh network to discover the failure and adjust the routing such that the data packets will reach the owner gateway.

In this experiment we started a TCP connection between *Client* and *Sky* and then moved the client in the vicinity of a different Internet gateway, forcing a gateway handoff to occur. Then we unplugged the power of the current Internet gateway. Figure 13 presents the evolution of a TCP

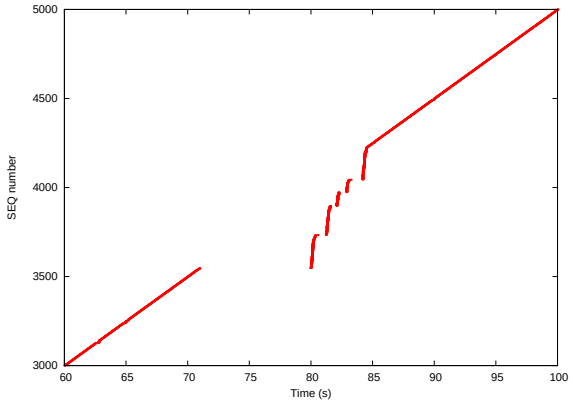

Figure 13: Inter-domain TCP fail-over test. Sky is the receiver.

flow where the X axis shows the time and the Y axis shows the packet sequence number. The graph starts after the first handoff from the original gateway. The graph shows about 8 seconds of disconnection required for the mesh network to detect the failure and adjust its routing. After that, it takes a few more seconds for TCP to catch up with the original rate. The network reacting to the failure in a timely manner prevented the disconnection of the TCP connection, overcoming the current Internet gateway crash.

Transmission Overhead: In this experiment we quantify the transmission overhead on the wireless and wired network when several mobile clients send and receive data. We performed experiments with varying number of clients, from 1 to 4, moving randomly inside the network. The cost for *maintaining the topology*, determined by the link state routing algorithm, was about 300 bps per mesh neighbor, independent on the traffic or the number of clients. Another component of the overhead is the cost of maintaining two multicast groups per client. The membership of these groups changes as the clients move and different access points join or leave the groups. The overhead generated by *managing multicast groups* depends only on the number of clients and is independent on the amount of traffic and/or flows in the network. In our topology this traffic amounts to about 500 bps at each access point per moving client. The traffic generated for *managing the clients* consists of Control Group traffic for access points coordination (about 3.7 Kbps per client when all clients were in the same vicinity), and the overhead generated for probing the link quality with the clients, which depends on the technique used (about 3.5 Kbps per client if DHCP requests are sent every 2 seconds, or negligible if RSSI is used but requires access points to support 802.11 monitor mode or driver support). Note that client management traffic exists only in the vicinity of the client and is not dependent on the amount of data traffic.

Internet gateways generate some overhead traffic on the wired network during the inter-domain handoff. Data pack-

ets are multicasted over the wired network to all other Internet gateways until the owner of the connection responds. In our tests, this process took between 10 ms and 80 ms. Note that data packets are forwarded in parallel to the end-host and their latency is much less. After the first handoff of a connection takes place, all Internet gateways are informed about the owner of that connection, and therefore new data packets are sent directly to the connection owner. As opposed to the wireless intra-domain overhead, which is only dependent on the number of clients, the inter-domain overhead is directly proportional to the number of connections each client has. However, the traffic generated by the inter-domain handoff is small, and uses only the wired connectivity.

6. Conclusion

In this paper we presented an architecture and an inter-domain routing protocol for multi-homed wireless mesh networks that provide uninterrupted connectivity and fast handoff. Our approach uses an overlay mechanism to integrate wireless and wired connectivity. We showed how overlay multicast groups are used to coordinate decisions between Internet connected access points to seamlessly transfer connections as the mobile clients move. The protocol optimizes the use of the wireless medium by short-cutting wireless hops through wired connections, paying a very low overhead during handoffs. We demonstrated the efficiency of our protocols through live experiments using an actual complete and available system, showing that the inter-domain handoffs occur instantaneously for both TCP and UDP connections.

References

[1] OpenWrt. http://openwrt.org.

[2] The Spines Overlay Network. http://www.spines.org.

[3] Akyildiz, I.F.; Jiang Xie; Mohanty, S. A survey of mobility management in next-generation all-ip-based wireless systems. *Wireless Communications, IEEE*, 11(4pp):16–28, Aug 2004.

[4] Akyildiz, I.F.; Wang, X; and Wang, W. Wireless mesh networks: A survey. *Computer Networks Journal (Elsevier)*, Mar 2005.

[5] Y. Amir, C. Danilov, M. Hilsdale, R. Musaloiu-Elefteri, and N. Rivera. Fast Handoff for Seamless Wireless Mesh Networks. *ACM MobiSys*, June 2006.

[6] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient overlay networks. In *Proc. of the 18th Symposium on Operating Systems Principles*, pages 131–145, Oct. 2001.

[7] M. M. Buddhikot, A. Hari, K. Singh, and S. Miller. Mobilenat: A new technique for mobility across heterogeneous address spaces. *MONET*, 10(3):289–302, 2005.

[8] D. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. *In Proceedings of MOBICOM 2003, San Diego*, 2003.

[9] K. M. H. Soliman, C. Castelluccia and L. Bellier. Hierarchical mobile ipv6 mobility management (hmipv6). June 2004.

[10] R. Hsieh, Z. G. Zhou, and A. Seneviratne. S-MIP: A seamless handoff architecture for mobile IP. In *INFOCOM*, 2003.

[11] K. Ramachandran, M. M. Buddhikot, G. Chandranmenon, S. Miller, K. Almeroth, E. Belding-Royer. On the design and implementation of infrastructure mesh networks. *WiMesh*, 2005.

[12] J. Moy. Multicast extensions to OSPF. RFC 1584, IETF, Mar. 1994.

[13] C. Perkins. IP Mobility Support. *RFC2002*, Oct 1996.

[14] J. Touch and S. Hotz. The x-bone. In *Third Global Internet Mini-Conference at Globecom '98*, Nov. 1998.

[15] V. Navda, A. Kashyap, S. Das. Design and evaluation of imesh: an infrastructure-mode wireless mesh network. In *6th IEEE WoWMoM Symposium*, June 2005.