

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

Yair Amir, Claudiu Danilov, Raluca Musăloiu-Elefteri, Nilo Rivera
{yairamir, claudiu, ralucam, nrivera}@dsn.jhu.edu

Department of Computer Science
Johns Hopkins University
Technical Report CNDS-2006-1 - July 2006
<http://www.dsn.jhu.edu>

Abstract—This paper presents a 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.

I. 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.

As a mobile device moves, a handoff between access points is required at two levels. An intra-domain handoff transfers the connectivity between access points to which the mobile device is directly connected. At a higher level, an inter-domain handoff between Internet gateways is required when the client connects to an access point that is closer to a different Internet gateway. While solutions exist for intra-domain handoff [1], [2], [3], we believe that currently there are no efficient and transparent protocols for inter-domain handoff in multi-homed wireless mesh networks. The routing protocol presented in this paper addresses the **inter-domain** handoff, providing fast handoff, and seamless connectivity as mobile devices move between Internet gateways located in different network domains.

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, extending our SMesh system [1] 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. Fast intra-domain handoff in SMesh is achieved by using multicast groups joined by access points in the vicinity of each client. These multicast groups are used for local coordination between access points that can potentially serve the mobile device, and for forwarding data packets to the client through multiple paths in periods of instability, until a designated access point is elected, thus providing uninterrupted service.

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 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 II presents related work. In Section III we describe the architecture of our multi-homed mesh network approach, and in Section IV we present the real-time handoff protocol between the wired Internet gateways. Experimental results are presented in Section V, and Section VI concludes the paper.

II. RELATED WORK

We propose a mechanism for multi-homed wireless mesh networks that optimizes routing and provide 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 [5]. 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) [6] 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 [8], [9]. 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.

Many reactive approaches have been proposed to address Internet connectivity in wireless ad-hoc networks [10], [11], [12], [13], [14]. A hybrid approach that achieves the same connectivity as in pro-active protocols but with less overhead was proposed in [15]. These schemes usually share similarities with MIP and although they are suitable for ad-hoc networks, they do not perform well in wireless mesh networks. Backbone nodes in a mesh network are stationary, as opposed to the nodes in ad-hoc networks, leaving space to more efficient protocols that exploit the relative stability of the mesh nodes.

Existing experimental wireless mesh testbeds that support client mobility include MeshCluster [2] and iMesh [3], both of which work with mobile clients in infrastructure mode. MeshCluster, which uses MIP for intra-domain handoff, shows a latency of about 700ms 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. Without actual layer-2 handoff latency accounted for (no moving client) and using layer-2 handoff triggers, layer-3 handoff takes 50-100ms. SMesh [1] provides 802.11 link-layer and network-layer fast handoff by working in ad-hoc 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.

III. 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 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




Fig. 1. Hybrid Overlay Mesh Network

gateways). There are several options available for a practical software router, including X-Bone [16], Spines [17] and RON [18]. In our implementation we use 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, access points located in different wireless islands can use their closest Internet gateways to relay packets between the islands.

A. Topology Formation: One Overlay Network

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, on which they periodically advertise their wired interface IP address. The multicast routing is handled by the underlying overlay infrastructure. 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.

B. 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 [19] as the wireless *ActualCost* metric, and latency as the wired links metric.

C. Hybrid Mesh Communication

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.

Peer Communication: Mobile clients advertise their presence periodically by broadcasting DHCP requests to renew their leases. Access points that are located in the vicinity of a mobile client, and that receive the client's DHCP requests, 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 client link quality and decide on which one of them should serve the client. 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.

Packets sent by a client to another 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




Fig. 2. TCP forward handoff: (a) Initial connection establishment (b) Connection handoff (c) Handoff completed

reason for using a multicast group instead of a single access point 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). It provides complete transparency for mobile clients, and a smooth intra-domain handoff between local access points. More details about intra-domain handoff in SMesh can be found in [1].

Client to Internet Communication: All Internet gateways join an anycast group using the overlay infrastructure. The anycast paths are created and maintained by the underlying multi-hop messaging system. Upon receiving a packet that has a destination outside the wireless mesh network, an access point simply forwards it to the anycast group joined by the Internet gateways. This way, packets are always sent to the closest Internet gateway. On the reverse path, a mechanism similar to the peer communication is used.

IV. INTERNET GATEWAY CONNECTION HANDOFF

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. These multicast trees may use Internet gateways (based on the routing metric) to either shortcut several wireless hops in the same island, or to hop between different wireless islands. 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.

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 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. However, this technique tends to be wasteful, as not all mobile clients may move and change their Internet gateway. In addition, 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.

In our approach, we treat UDP and TCP connections separately, detect the existing owner (the Internet gateway from which the connection was initiated) of a connection on the fly, and forward existing connections through their original owners, while new connections are relayed through the closest gateway. Below, we describe our inter-domain handoff mechanisms.

A. 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. When an Internet gateway receives a TCP packet that is not a SYN, if the gateway does

not have an entry for that connection in its NAT table, it forwards that packet to the Internet gateway multicast 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 Internet gateway multicast 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 handover process. Figure 2 shows the stages of such a TCP connection handoff.

If packets arrive at an Internet gateway at a fast rate, there may be several packets sent to the Internet gateway multicast 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).

B. 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 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. Connection-less 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

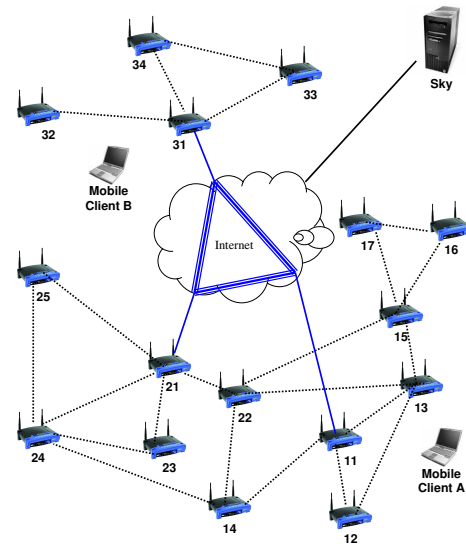


Fig. 3. The Hybrid Testbed.

certain timeout (in our implementation 1 second), the Internet gateway will claim ownership of the UDP connection, will stop forwarding packets of that connection to the Internet gateway multicast group, and will continue to relay packets to the Internet.

Due to handoff 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 both gateways claiming the ownership of the connection. In such event, upon receiving a response from a gateway that claimed ownership already, the gateway with the highest IP gives up its ownership.

V. EXPERIMENTAL RESULTS

A. Setup

We implemented our protocols within the firmware of Linksys WRT54G wireless routers. A third party firmware (EWRT [20]) 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 3. 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




Fig. 4. P2P UDP test. Latency of packets received at Client A.




Fig. 5. P2P UDP Test. Latency of packets received at Client B.




Fig. 6. P2P UDP Test. Lost packets at Client A.




Fig. 7. P2P UDP test. Lost packets at Client B.

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 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 end 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 3rd 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 handoff 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 20ms at a rate of 64Kbps, for 5 minutes. We focus our experiments on VoIP as a representative application that poses severe latency requirements.

B. Measurements

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 4 - 7 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 4 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 4 and 5 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 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 6 and 7 present the packets lost at the two clients during the experiments. Loss is represented as cumulative number of losses over the last 25 packets. For example, Figure 7 shows the highest cumulative loss around packet 2000, where 15 consecutive packets were lost. Note that losses are not related to access point handoffs. 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 8 and 9 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 $4ms$ 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 handoff, a better connected access point starts serving the client and the packet latency decreases.

Figures 10 and 11 show the packets lost at *Client* and *Sky*, respectively, in the same experiment. We can see that losses occurred in bursts of less than 7 packets in both streams. The number of packets that arrived after more than $100ms$ 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 (15,000 in each direction), very few packets were lost or delayed.

Figure 12 shows the duplicate packets from *Sky* to *Client*. These duplicate packets appear during intra-domain handoffs, and occur because multiple access points in the vicinity of the

mobile client join the client's Data Group. As they all forward the data packets, duplicates are received at the mobile client. As soon as the access points in the vicinity of the client agree on which one of them should serve the client, the other access points leave the Data Group and duplicates stop.

In Figure 13 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 Internet gateway multicast group and to the final destination (as explained in Section IV-B). 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 Internet gateway multicast 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: While our protocols offer real-time handoff for UDP-based applications, it is important to see how TCP streams are affected by the Internet gateway handoffs. Figure 14 shows the latency of a TCP stream from *Client* to *Sky*. We first note that the TCP connection did not break during handoffs. Between packets 5500 and 6000, we see the latency jump up and down by about $4ms$ due to either using the original Internet gateway that owns the connection, or using a different gateway that forwards the connection to the original access point owner. The total number of packets that arrived after $100ms$ was 243, compared to 92 in the connection-oriented UDP test. This is because TCP delays packets that arrive out of order (mainly due to lost packets in our case) 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 15 presents the evolution of a TCP 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




Fig. 8. Client-Internet UDP test. Client is the receiver. Packet latency.




Fig. 9. Client-Internet UDP test. Sky is the receiver. Packet latency.




Fig. 10. Client-Internet UDP test. Client is the receiver. Lost packets.




Fig. 11. Client-Internet UDP test. Sky is the receiver. Lost packets.




Fig. 12. Client-Internet UDP test. Client is the receiver. Duplicate packets.




Fig. 13. Client-Internet UDP test. Sky is the receiver. Duplicate packets.

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. This includes the costs of maintaining the topology, managing the

multicast groups (access point joins and leaves), managing the clients and performing the intra and inter-domain handoff. We performed experiments with varying number of clients, from 1 to 4, moving randomly inside the network. For each access point, we recorded the following types of traffic:

- *Maintaining the topology:* A hello message is exchanged between directly connected access points every 2 seconds. For each node, this amounts for about 300bps per mesh neighbor, independent on the number of clients or data traffic. The traffic generated by the link state routing




Fig. 14. Client-Internet TCP test. Sky is the receiver. Latency graph.




Fig. 15. Client-Internet TCP fail-over test. Sky is the receiver.

protocol was negligible as the topology was stable during the experiments (the mesh access points are stationary).

- *Managing the multicast groups:* We assign two multicast groups per client. The membership of these groups changes as the clients move and different access points join or leave the groups. In our topology this traffic amounts for about 500bps per moving client ($\sim 300\text{bps}$ sent and $\sim 200\text{bps}$ received at each access point).
- *Managing the clients:* Each client advertises its presence using periodical DHCP requests that amount to about $\sim 3.5\text{kbps}$ per client. The traffic generated on the Control groups for access point coordination was up to 11kbps for 4 clients, when all clients were in the same vicinity. The DHCP requests and the coordination traffic multicasted to the control group are sent only in the vicinity of the client. They are not forwarded to other nodes in the mesh network.

Internet gateways generate additional traffic on the wired network due to inter-domain handoff. Data packets 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 10ms and 80ms . After the first handoff of a connection takes place, all other Internet gateways are informed about the owner of that connection, and therefore new data packets are not multicasted, but rather sent directly to the connection owner. As opposed to the wireless intra-domain overhead, which is mainly 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. Overall, the overhead of the system is quite small compared with 802.11 thoughtputs (tens of Mbps).

VI. CONCLUSION

In this paper we presented 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 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] Yair Amir, Claudiu Danilov, Michale Hilsdale, Raluca Musaloiu-Elefteri, Nilo Rivera, "Fast Handoff for Seamless Wireless Mesh Networks," *ACM MobiSys*, June 2006.
- [2] G. Chandranmenon S. Miller K. Almeroth E. Belding-Royer K. Ramachandran, M. M. Buddhikot, "On the design and implementation of infrastructure mesh networks," *WiMesh*, 2005.
- [3] V. Navda, A. Kashyap, S. Das, "Design and evaluation of imesh: an infrastructure-mode wireless mesh network," in *6th IEEE WoWMoM Symposium*, June 2005.
- [4] Wang X. Akyildiz, I.F. and W. Wang, "Wireless mesh networks: A survey," *Computer Networks Journal (Elsevier)*, Mar 2005.
- [5] S. Akyildiz, I.F.; Jiang Xie; Mohanty, "A survey of mobility management in next-generation all-ip-based wireless systems," *Wireless Communications, IEEE*, vol. 11, no. 4pp, pp. 16–28, Aug 2004.
- [6] C. Perkins, "IP Mobility Support," *RFC2002*, Oct 1996.
- [7] Milind M. Buddhikot, Adishesu Hari, Kundan Singh, and Scott Miller, "Mobilenat: A new technique for mobility across heterogeneous address spaces," *MONET*, vol. 10, no. 3, pp. 289–302, 2005.
- [8] K. Malki H. Soliman, C. Castelluccia and L. Bellier, "Hierarchical mobile ipv6 mobility management (hmipv6)," June 2004.
- [9] Robert Hsieh, Zhe Guang Zhou, and Aruna Seneviratne, "S-MIP: A seamless handoff architecture for mobile IP," in *INFOCOM*, 2003.
- [10] Y. Sun, E. Belding-Royer, and C. Perkins, "Internet connectivity for ad hoc mobile networks," *International Journal of Wireless Information Networks*, 2002.
- [11] Yuan Sun and Elizabeth M. Belding-Royer, "Application-oriented routing in hybrid wireless," *ICC '03*, 2003.
- [12] W. Matthew, J. Miller, and N. Vaidya, "A hybrid network implementation to extend infrastructure reach," *UIUC Technical Report*, 2003.
- [13] Ulf Jönsson, Fredrik Alriksson, Tony Larsson, Per Johansson, and Gerald Q. Maguire Jr., "Mipmanet: mobile ip for mobile ad hoc networks," in *MobiHoc*, 2000, pp. 75–85.
- [14] Yu-Chee Tseng, Chia-Ching Shen, and Wen-Tsuen Chen, "Integrating mobile ip with ad hoc networks," *Computer*, vol. 36, no. 5, pp. 48–55, 2003.
- [15] R Kravet P Ratanchandani, "A hybrid approach to internet connectivity for mobile ad hoc networks," *IEEE Wireless Communications and Networking Conference*, 2003.

- [16] J. Touch and S. Hotz, "The x-bone," in *Third Global Internet Mini-Conference at Globecom '98*, Nov. 1998.
- [17] "The Spines Overlay Network," <http://www.spines.org>.
- [18] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient overlay networks," in *Proc. of the 18th Symposium on Operating Systems Principles*, Oct. 2001, pp. 131–145.
- [19] D. De 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.
- [20] "Portless Networks, EWRT," <http://www.portless.net/menu/ewrt>.