Routing and Mobility Performance in Wireless Base-Station Networks[1]

Stuart D. Milner, Karthikeyan Chandrashekar, Sohil Thakkar, and Chieh-Hung Chen

Clark School of Engineering
University of Maryland
College Park, MD 20742
E-mail: milner@isr.umd.edu; {karthikc,sohil,jaychen}@glue.umd.edu

ABSTRACT

This paper focuses on the scalability of mobile, wireless, base-station-oriented networks. In our study of base-station networks, we investigate the performance limits and communication complexity of the protocols and services under mobility and subject to dynamic behavior of hosts and base-stations. In addition, we discuss the requisite development and implementation of the base station network architecture made to the wireless models in OPNET version 7. These include: multiple wireless interfaces, an exclusive control channel, a mobility "handoff" process, a "semi-broadcast" OSPF routing algorithm, and router (IP) mobility functions.

INTRODUCTION

We are investigating the scalability of wireless networks, dynamic base-stations and hosts. The work is based on, and enhances, the architecture and development of the DARPA-funded On Board Switch (OBS) Project [1]. For dynamic base-station-oriented networks ([2], [3]), we are investigating the scalability limits of both routing and mobility management protocols and their associated services. In this approach, the routers in the base stations are dynamic and hosts move within and between IP router domains. We are considering networks with as many as thousands of dynamic users, links and routers, and our simulation approach is based on a complete systems view of scalability including physical, link, network and transport layer models and associated processes.

Base-station-oriented architectures, in contrast to ad hoc networking architectures, are inherently hierarchical, with mobile hosts connected to "advantaged" base station nodes (routers, switches, high data rate transmitters and receivers, high speed encryption devices, etc.) [4]. Such nodes may be satellite-, airborne- and/or terrestrially based, and their topology (links and switches) is dynamic. They constitute a wireless intranet, in effect, but can fully interoperate with the larger Internet. Moreover, such networks are rapidly deployable, "instant infrastructures" that are comprised of large host to router ratios [3]. Routing and mobility management occurs within the mobile domain or network subset, and mobile users move within and between the mobile base stations — in a cellular-like fashion. Such networks are hierarchical, and, as such, hold promise for scalability.

ARCHITECTURE AND SCALABILITY

The OBS project, and the Warfighter's Internet [3] provided for a unified wireless infrastructure, which included the convergence of internet (packet switching, quality of service, routing, transport), cellular telephony (mobility, handoffs, base stations), and wireless multiple access (802.11, dynamic TDMA) technologies [5,6]. The underlying architecture includes mobile hosts and base stations (routers) forming a network as seen in Figure 1. The base station nodes, which consist of routing, mobility and radio functions, form a wireless, high data rate backbone. Mobile hosts can affiliate with any router through a handoff process.

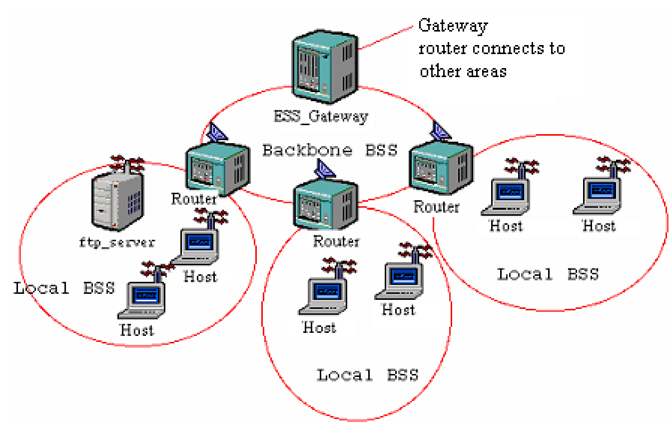

Figure 1 — Base Station Architecture

The scalability of the Internet is due in large part to its hierarchical nature. Likewise, the hierarchical nature of this base-station oriented wireless networking architecture is the major factor contributing to its scalability (Figure 1). In this

[1] This work was supported in part by the Defense Advanced Research Projects Agency.

all wireless implementation, a set of hosts and their associated router form a LAN. All hosts directly attach to a router, in essentially one hop. At the next level, a set of routers forms a backbone LAN, and at a higher level, a set of gateway routers form a backbone linking OSPF areas or router subnets. In all of these cases, the wireless LANs uses the IEEE 802.11 standard.

MODIFICATIONS TO OPNET WIRELESS LAN MODEL

Modeling and simulation of this network, where all base stations/routers and hosts can be mobile and all levels of the hierarchy are wireless required that the routers have at least two wireless interfaces, one for the host/router LAN and one for the backbone LAN. Initially, the wireless model of Version 7 allowed only one wireless and one wire line interface per router. Accordingly, we added additional wireless interface capability to the node model of the standard wireless LAN router. Furthermore, to support routing over a wireless backbone requires broadcast packets (*Hello* messages in the case of OSPF) to be transmitted over the wireless interface. The standard OPNET model did not allow broadcast of packets over a wireless interface. As a result, modifications to the MAC process were made to enable broadcast over the wireless backbone.

Our model requires that the hosts be able to move between 802.11 LANs or Basic Service Sets (BSS). In the standard OPNET model each subnet represents one BSS, and a node cannot be assigned a trajectory that will allow it to cross subnets. To allow hosts to move between BSSs, we introduced the notion of a BSS id**,** which determines the current subnet of a node. Thus the OPNET "subnet" is now a logical subnet and can contain many BSSs within it. Each BSS in the subnet has a unique frequency assignment, and this frequency assignment procedure is automatic and a function of the BSS id. When a node is initialized, its BSS id will determine the unique BSS or subnet association.

In the process of validating our models and simulating large number of nodes in heavy load scenarios, we found that the OPNET wireless 802.11model was failing due to the large number of interrupts being generated. Each occurrence of a statistical interrupt in the back-off state of the MAC process diminished the back-off slot time until this value becomes negative. This resulted in erroneous interrupts scheduled. The problem was fixed by changing the condition under which the expression for the back-off slot will be evaluated.

All of the above modifications have been incorporated in the standard wireless models of OPNET 8.0.

MOBILITY ARCHITECTURE

The mobility protocol for hosts is similar in concept to the standard Mobile IP protocol. In addition, the protocol allows a mobile host to bind with the router providing the best service (e.g., maximum signal strength; least congestion) and, in the process of binding, to maintain its IP address. Following key design features describe the protocol [7,8]:

- All router nodes will periodically advertise a set of preference values to indicate the status of the router node.
- A mobile host, which does not receive an advertisement from the router to which it is presently bound, will send a solicitation message to the router to explicitly request an advertisement.
- A mobile host binds to the router node providing the best service as determined by the preference value (e.g., status of the router; congestion) and signal strength measurements. The host evaluates its binding status with respect to the latest set of preference advertisements, and may affiliate to a new router node by initiating a local binding.
- The mobile host after acquiring a local binding with a new router may perform a remote binding to bind to its home router node. The home router node is then capable of forwarding all packets addressed to the mobile host to the address of the router node with which the host is presently affiliated.

Mobility Protocol Design and Implementation

All mobility messages are implemented as MAC layer messages, and the mobility processes are implemented above the MAC layer in the host and the router node models. All mobility messages are sent over a separate control channel**.** The OPNET node models for the host and router were modified to incorporate a fifth control channel. All mobility control messages are sent over the global control (i.e., fifth) channel, except for the remote binding message. The fifth control channel thus provides an additional capacity of 11 Mbps over the existing four data channels. All nodes within hearing range (wireless LAN range for reception of packets without interference, the default setting in the OPNET model is 300 meters) of the sender can receive messages on the control channel. The purpose of this additional channel is to enhance the scalability of the network.

The control channel, like the data channel, supports reliable data delivery. All mobility messages over the control channel are treated as equivalent to data messages and, therefore need to be acknowledged. The mobility advertisements, however, are broadcast packets, and according to the protocol are not acknowledged.

Each host is initially bound to the router in its BSS. The host periodically evaluates its current binding and may decide to bind to another router, based on the mobility advertisement information broadcast by the routers. The period of these router advertisements is a configurable parameter. The advertisements contain the preference value of the router; this value indicates the status of the router and is used by the hosts to evaluate their binding decisions. After receiving the first advertisement, the host waits for a configurable interval, to ensure that it has received sufficient number of advertisements before making the handoff decision. The Host then selects the three best routers (signal strength to the host) and among these chooses the one with the highest preference value. If the host hasn't received an advertisement from the router to which it is bound, it will send a solicitation message to its router to request for an advertisement before making the binding decision.

In order to prevent the host from making frequent handoffs for marginal increase in signal performance, each host is assigned a binding threshold. This configurable parameter enables the host to evaluate its current binding. If the service provided by the router to which the host is presently affiliated is providing a better than the binding threshold, then the host will not hand off to another router even if the new router is providing better service. The host on deciding to affiliate to a new router initiates a local binding procedure. Once the local binding is successful the host may initiate a remote binding procedure to its home router (the router that has the same subnet mask as the host).

Host Initiated Advertisements

Based on our experimentation, we determined that loss of advertisement packets sent by the home router to its locally attached host resulted in additional bindings. This is because the host is forced to bind to a new router. However, this could not be the correct decision as the loss of the advertisement packet does from the current router could simply mean that the packet is dropped due to collision or is corrupted due to interference. To prevent such incorrect binding decisions the host will send a solicitation message to the current home router if its does not receive the mobility advertisement from the current router. The home router upon receiving the solicitation message from the host, waits for a small interval of time to collect solicitations from other hosts and then re-broadcasts the advertisement with the same preference set. The host waits for the advertisement for a finite period of time before making the binding decision.

Local Binding procedure

Once the handoff decision is made and a router has been chosen, the host will send a local bind request message to the new router node on the global control channel. The router on receiving the message includes the new host in its BSS by adding the host MAC address to the MAC address list of the BSS. The IP address of the new host is also stored in a local bind list, which will be used by the IP process while forwarding packets to the locally bound host. The router sends a local bind reply as a confirmation of the binding process.

The host**,** on receiving a reply, will change to the new routers BSS by changing its BSS id. The host will also modify its frequency (which is BSS based) so that it can receive and send messages in the new 802.11 BSS. It will then record the MAC address of the new router as its AP (Access Point) and the IP address of the new router as its default gateway. After making these changes the host will release its previous local binding by issuing a local bind release message to the previous router.

The previous router after receiving the release message will modify its local bind list and the BSS based MAC address list**.** This will ensure that these lists will reflect the current status of the network and will prevent any looping of data (the lists are used by IP for data forwarding, outdated entries may result in an incorrect path or a loop).

Once a host has established a local binding with a new router, the new host will be treated as if it belongs to the new routers BSS even though it has a different IP address. The modifications made to the IP and MAC processes will ensure that all data generated within the BSS for the new host and all data generated by the new host to any host in the new router BSS will be correctly forwarded. The IP process will use the entries in the local bind list to forward packets within the same BSS.

Remote Binding Procedure

The host after binding locally to a new router (any router other than its home router) will initiate a remote binding with its home router by sending a remote binding request.

The remote binding message is transmitted as an IP datagram to ensure that the home router receives the remote binding message even if it is beyond the 802.11 communication range. All packets to the mobile host are forwarded through the home router (except for source nodes belonging to the same BSS as the one recently joined by this mobile host). Hence, it is imperative that the home router learns the current location of any host belonging to its BSS.

This remote binding request message contains the IP and MAC address of the host and the default gateway IP address of the new router (care of address of the host). The home router then updates its remote bind list to reflect the new forwarding address for that host. The IP process will use the remote bind list to forward packets addressed to the host that is now remotely bound to it.

Once a remote binding is established, the modifications to the IP and MAC processes will ensure that the home router will forward the packets designated for a host that belongs to its subnet. The home router will forward the packets to the care of address of the host. The care of address of the host is its foreign router, which will then forward the packets to the destination host using the host entry in the local bind list.

SCALABILITY OF THE MOBILITY PROTOCOLS

By design, all mobility messages except the remote binding message are transmitted via the exclusive mobility or fifth channel. Hence, the mobility overhead is isolated and will not affect the data channel. The two factors that determine the mobility channel overhead are: 1) number and the interval of mobility advertisements; and, 2) the number of bindings (local and remote).

The number of routers in the network and their advertisement intervals determine respective traffic load and throughput. To study the load pattern of the advertisements a network with 20 routers and a single host within range of all routers was simulated to generate mobility traffic for a 30-minute period. Binding capability in the host was disabled to ensure that the data on the control channel represents mobility advertisements only.

The relationship between advertisements, number of routers, offered load (data rate), and throughput (data received) is given in Figure 2. Both load and throughput (in bps) are given for different advertisement intervals, and increase in a linear manner with the number of routers. For 20 routers with the advertisement interval set to 15 seconds, the load is 225 bps averaged across the 30 minute simulation. The load due to the advertisements is less than **.01% of the effective capacity of the 11Mbps channel** (effective capacity defined as 30% of 11Mbps).

Furthermore, as can be seen in Figure 2, throughput decreases relative to the offered load as the number of routers increases – the divergence indicating that loss of advertisement messages is due to collisions. The throughput equals the offered load in the case of 5 routers and is ideal regardless of the advertisement interval. Adding random jitter to the advertisement intervals can reduce the number of collisions for the larger number of routers. This, however, will increase the total time required for the binding process to complete, as the host will now have to wait for a longer duration to collect sufficient advertisements to make the binding decision.

Clearly, the load due to advertisements is negligible and the number of advertisement packets lost due to collision is small. For example when advertisement interval is 15 seconds, **with 10 routers, 1.5% of the total advertisement packets are lost and with 20 routers, 3% of the advertisement packets are lost**.

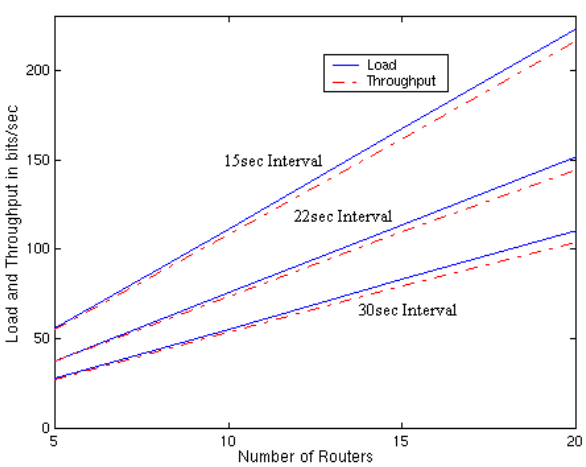

Figure 2. Load and throughput due to mobility advertisements

To study the load in the mobility channel, a network with 55 routers distributed over 5 areas, including 1 gateway router per area, and 10 hosts affiliated with each router was simulated for 30 minutes of real time. In this experiment, the routers in each area are fixed and are arranged in a horizontal line. The hosts are mobile and the trajectory assigned to them will cause the hosts to move towards the right most router. After each advertisement interval, the routers broadcast their preference sets and all hosts subsequently make the binding decisions based on these preference sets and signal strength measurements.

Figure 3 shows the relationship between load and the number of bindings. As can be seen this relationship is linear. The peak load in the channel is approximately 123 Kbps, with 340 of the 500 hosts simultaneously exchanging binding messages. The binding messages as well as the preference advertisements constitute the mobility channel load, and the **total load is less than 3.5% of the effective capacity (30%) of the 11 Mbps mobility channel.**

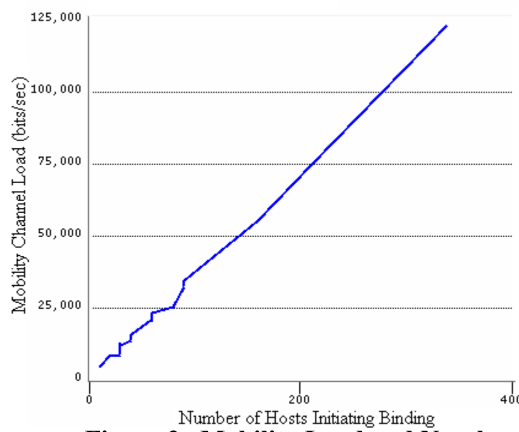

Figure 3. Mobility Load and Number of Bindings

Using the same scenario, the success rate of the bindings was also studied (Figure 4). A successful binding is one for which the concerned host receives the local bind response message, in effect completes the local binding procedure. As can be seen all the hosts initiating a binding succeeded. This is due to the effectiveness of isolating traffic using the control channel.

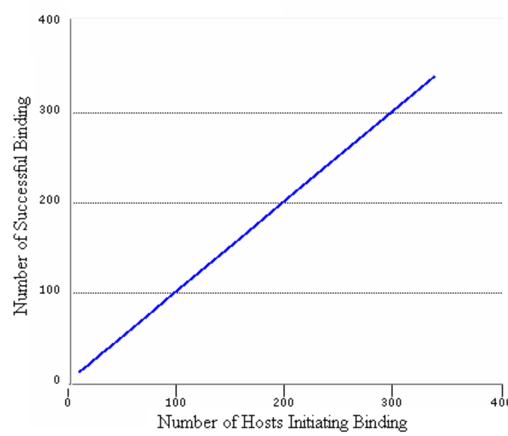

Figure 4. Binding success rate

BASE-STATION ROUTING APPROACH

The routing protocol used in this architecture as well as in the OBS project is Open Shortest Path First (OSPF) [7]. Certain enhancements were made to OSPF including flooding of topology information, predictive routing and the incorporation of "semi-broadcast links" [8]. The latter modification was made in order to customize OSPF for use in radio networks, where one-hop reach ability between router nodes cannot be guaranteed; hence the need for a single Designated Router (DR) used in standard versions of OSPF was eliminated. To compensate for the lack of a centralized routing database, each router has to exchange topology information with every other router it can directly reach to form an adjacency; thereby creating dynamic point-to-point links between all routers in a range. With this capability, each router is represented as a vertex and the dynamic link as an edge on a topology graph.

If this were implemented using a wire line version of OSPF, the communication complexity would be of the order $O\ (N\times E)$, where N is the number of routers and E is the number of links [9]. Since wireless transmissions can often reach multiple destinations, it is hypothesized that communication complexity in the wireless model should not exceed the wire line model. This increase in complexity can be mitigated by localizing router traffic into areas (Extended Service Sets in 802.11 terms; OSPF "areas" in IP terms).

Modifications to OPNET OSPF Model

We have added new Point to Multipoint interface type to the OPNET model of OSPF. This interface type support dynamic multiple point to point link between routers, which are in range. With this capability, each router is represented as vertex and dynamic point-to-point link as an edge on a graph.

The scalability of the routing is based on a new OSPF link type ("semi-broadcast") that enables area wide topology exchange and routing, and by connecting different area in hierarchical fashion [7,8]. The net effect is to maximize transmission of overhead packets in respective "areas" or ESSs.

SCALABILITY OF THE ROUTING PROTOCOL

In order to simulate mobility conditions for the routers, we established an initial topology of 20 routers configured in a linear path; the distance within the routers was set at 250 meters. Based on maximum communication range limits that we set, only adjacent routers could communicate with each other. Then, in successive 3-minute intervals and from left to right, the routers moved until there was an overlap in physical location, creating new adjacencies each time. The effect of this mobility scenario was to change and create new adjacencies, and, thereby allows analysis of the respective OSPF traffic overhead [Hello and Link State Update (LSU) messages].

The routing update time as it relates to the number of adjacency (router) changes is given in Figure 5. The intervals associated with link state updates due to adjacency changes and router movement range from 7 to 28 seconds. There appears to be no relationship between number of adjacency changes and associated update interval. In a reverse experiment when the adjacencies were decreased from the maximum of 16 to 2.Similar observations were noted for experiments involving 5 areas (55 routers including 5 gateway routers). Therefore, the OSPF update interval time does not depend on number of adjacency changes.

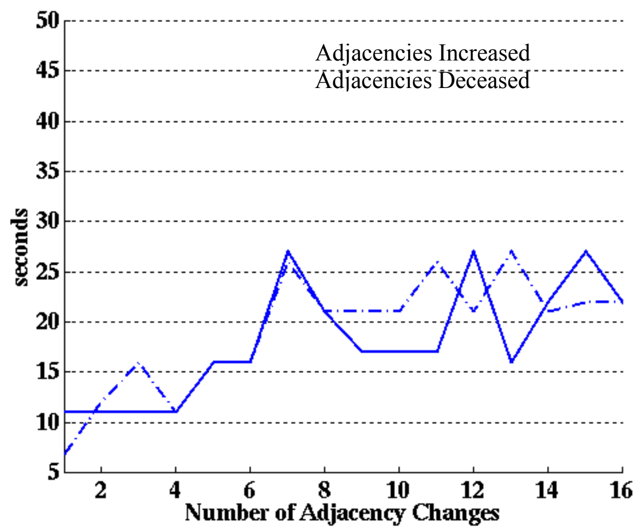

Figure 5. Routing Update Time

We noted that traffic pattern for this routing update interval is negatively skewed and contains a large peak of one-second duration. The aggregated peak data rate increases as the number of adjacency changes increase due to the flooding of routing update information. In order to determine the effects of subnetting the routers, we conducted experiments comparing the number of adjacency changes for 20 routers distributed in one area with the same number of adjacency changes, equally distributed in two areas (Table 1). In the case of 20 routers in one area, the aggregated peak data rate varied from 108 kbps with two adjacency changes to 457 kbps with eight adjacency changes. When we divided the routers into two subnets, we noted a significant decrease in the aggregated peak data rates due to link state exchanges for the same number of adjacencies (30% reduction with eight changes). This suggests that the number of routers within an area should be kept small while increasing the number of areas. It also suggests that by increasing subnets, flooding is reduced, hierarchy is achieved, and scalability is facilitated. The optimal number of routers in a subnet is subject to further research.

	Number of Adjacency changes			
	2	4	6	8
Case I	**108**	**215**	**311**	**457**
Case II	**22**	**39**	**88**	**136**

Note: Case I – 20 routers in 1 area
Case II – 20 routers in 2 areas (10 each area)

Table 1. Aggregated peak data rate (kbps) due to number of adjacency changes

Figure 6 below shows the total number of messages transmitted as a function of adjacency changes. As the number of adjacency changes increases, the number of OSPF message exchanges increases. As in the previous analysis, similar observations were noted for experiments involving 5 areas (55 routers including 5 gateway routers).

Moreover, the number of packets or messages is of the same order of magnitude for the adjacency changes. The observations on communication complexity support the earlier hypothesis that due to the semi-broadcast nature of wireless networks, a single transmission may be sufficient in exchanging link state information with adjacent routers; hence the communication complexity of the wireless network may not exceed that of a wire line network.

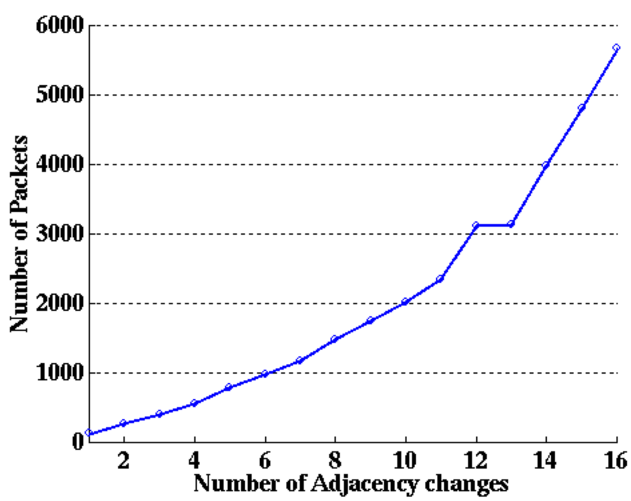

Figure 6. OSPF Message Complexity

Combined Performance of Routing and Mobility protocols

To study the combined impact of the routing and mobility protocols the following model was simulated. The model incorporated 5-OSPF-areas. Each area has 10 routers and 10 hosts affiliated with each router, and 5 FTP servers. A backbone LAN consists of 5 gateway routers, which has interfaces to both the respective area router backbone and to the gateway router backbone. Therefore, each experiment includes 500 hosts, 55 routers, and 25 FTP servers. However, the backbone of five gateway routers was implemented as an Ethernet. Nevertheless, it should have no real effect on the results as these gateway routers, in effect, are fixed wireless nodes.

Two scenarios were modeled. In the first scenario all the 580 nodes in the network were fixed. In the second case, all nodes except the gateway routers were assigned trajectories. The mobile routers will cause adjacency changes, and the mobile hosts will cause change in affiliation. Figure 7 below shows the offered load for all LAN traffic for both the fixed and all-mobile scenarios. The load in the mobile scenario ranged from 1.17Mbps to 8.58Mbps and is less than the fixed wireless case, which ranged from 2.35Mbps to 10.95Mbps. The decrease in the this load is due to the availability of alternate routes for packet transfer as a result of the occurrence of favorable adjacency changes In other words, the alternate routes will reduce the number of the hops the packet has to travel to reach the destination. Since in our model the number of adjacency changes is increased as the simulation proceeds, the size of the peaks increases. Therefore, the routing updates reduce the load in the network.

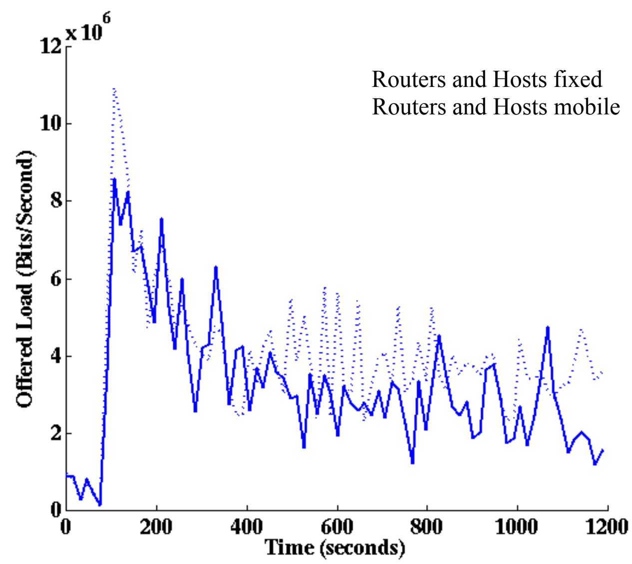

Figure 7. WLAN Offered Load

SUMMARY AND CONCLUSIONS

Our results confirm the scalability of mobile, wireless, base-station oriented networks. The design of routing and mobility protocols contributes to performance and complexity scalability.

In order to implement these models, we enhanced OPNET software to support: multiple wireless interfaces; multiple BSS based networks; a separate control channel for mobility messages; a back-off state in the wireless MAC process was modified for large-scale simulations. Most of our modifications were incorporated into OPNET, Version 8. In addition, we implemented an enhanced version of the mobility and routing protocols, specifically designed and developed for mobile, wireless, base-station networks. OPNET provides the capabilities to conduct simulations for unified, mobile wireless infrastructures.

We studied the overhead associated with the mobility protocol and found it to be minimal. The average load due mobility advertisements grew linearly as the number of routers within a communication range increased. Average load due to binding also increased linearly as the number of hosts initiating binding increases

Our investigation of routing overhead suggested that it depends on adjacency changes and total number of routers per area but not on application traffic. The routing overhead was significantly greater than the mobility overhead.

We found that the number of routers within an area or extended service set should be kept small while increasing the number of areas. We found that by increasing areas or subnets, flooding is reduced and scalability is facilitated.

In a large-scale experiment with 580 router, FTP server and host nodes, the aggregated average wireless LAN load ranged from approximately 2.35 Mbps to 10.95 Mbps. The offered load when all nodes were mobile was less, in general, than under fixed wireless conditions.

Our investigations focused on the degree of mobility (i.e., change in router adjacencies and host bindings). The protocol overhead associated with this mobility is scalable.

We conclude that the mobility and routing architecture of base-station oriented networks is a key factor in the scalability of such systems. Furthermore, the hierarchical nature of the architecture facilitates scalability.

.

ACKNOWLEDGMENT

The authors acknowledge the contributions of Janice Chao, Wei-Lun Chen, and Tzung-Hsien Ho of the University of Maryland and the OPNET Model Research and Development group to this research project.

REFERENCES

[1] "Final Report for On Board Switch," BBN Technologies and TRW Tactical Systems, 22 January 2000.

[2] NRC, *The Evolution of Untethered Communications*, Washington, DC: National Research Council, 1997.

[3] S. Milner, S. Bernstein, C. Niessen, and D. White, "Wireless Warfighter's Internet," *Proceedings of MILCOM'97*, Volume 1, November 1997.

[4] C. Perkins, ed., *Ad Hoc Networking,* Boston: Addison-Wesley, 2001.

[5] OPNET, "Wireless LAN (IEEE 802.11) Model Guide," 2000.

[6] "Deployable and Adaptive Mobile Ad Hoc Networks." Hughes Research Laboratories and Sarnoff Corporation, prepared for the Defense Advanced Research Projects Agency, November 1999.

[7] "Software Design Document for On Board Switch," BBN Technologies and TRW Tactical Systems, 22 January 2000.

[8] "Protocol Design Document for On Board Switch," BBN Technologies and TRW Tactical Systems, 22 January 2000.

[9] J.F. Kurose and K.W. Ross. *Computer Networking: A Top-Down Approach Featuring the Internet*, Addison Wesley, Reading, MA (2000)