Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

SIMULATION-BASED ROUTING PROTOCOL PERFORMANCE ANALYSIS – A CASE STUDY

Talal M. Jaafar
George F. Riley
Dheeraj Reddy

Department of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332, U.S.A

Dana Blair

Cisco Systems
500 Northridge Road, Suite 800
Atlanta, GA 30350, U.S.A

ABSTRACT

In this paper we evaluate the Enhanced Interior Gateway Routing Protocol *(EIGRP)* via packet simulations. EIGRP, an intra-domain routing protocol developed by Cisco, is mainly based on the Diffusing Update Algorithm *(DUAL)* which computes shortest paths distributedly without creating routing-table loops or incurring counting-to-infinity problem. Previous studies showed EIGRP's ability to adapt quickly to routing changes in medium-scale networks. In our research, we developed a detailed simulation model of EIGRP (publicly available), and we used it to evaluate EIGRP performance under a very dynamic network. Our results showed that EIGRP converges faster than a single TCP timeout in most cases. The simulated network was a composite of wired and wireless hosts, and the results hold for both types of media. In addition, the study showed a feasible approach for seamless mobility and continuous connectivity for users of mobile wireless devices as they move within an Autonomous System (AS).

1 INTRODUCTION

Since its development, EIGRP is known to converge as quickly as a link-state protocol while being loop free. The EIGRP developers (Albrightson, Garcia-Luna-Aceves, and Boyle 1994) state that many medium-scale network studies were performed and EIGRP proved to be a robust and reliable intra-domain routing protocol. As of 2000, network architects (Pepelnjak 2000) state that EIGRP and Open Shortest Path First *(OSPF)* (Moy 1989) are being implemented in approximately half of the networks. EIGRP is not only an enterprise-oriented routing protocol, but also a protocol that can be used in service-provider environments because it has fewer topology limitations than OSPF (Pepelnjak 2000). However, OSPF has seen a wider spread in the service-provider market because most new service-provider oriented technologies such as MPLS/VPN (Virtual Private Networks based on Multi-Protocol Label Swapping) are first implemented within the framework of OSPF.

These routing protocols are sophisticated distributed algorithms and a deep understanding of their performance/behavior is very crucial and hard as they are deployed in medium to large-scale networks. Simulation tools have been typically used in computer network systems study to evaluate current/new architectures or perform systems tuning. However, there is usually a trade-off between accuracy and scalability. A full detailed model of the system will require a large amount of memory and cpu power, and sometimes unfeasible runtime. On the other hand, an abstracted model will result in achieving large-scale simulation as well as less characterization of the system. Although EIGRP is widely deployed, there have been limited studies around it. Opnet (Bertolotti and Dunand 1993) which is a communication network simulator has an EIGRP model that computer researchers could use; however, Opnet is not available publicly, and its simulations have limited scalability which in turn may not capture the main characteristics of EIGRP.

In our research, we provide a scalable and detailed simulation model of EIGRP to evaluate its performance. Our performance analysis of EIGRP is conducted on a small, but very dynamic network. The network is a mix of wired and wireless networks. We chose to simulate such network to provide a high rate of route changes. The route recomputations are not triggered by taking links/interfaces down or up, rather by the wireless handoffs of the mobile hosts within the network. We will explain later on in detail how this is accomplished. Our results show the ability of EIGRP to update routing tables in a timely fashion, usually within a single TCP timeout period in an AS. Moreover, the experimental results motivate a new approach for mobile computing which is discussed later in the paper.

The remainder of this paper is organized as follows. Section 2 discusses the simulation framework and our implementation efforts of EIGRP and wireless-handoff in Georgia Tech Network Simulator *(GTNetS)* (Riley 2003) to support

1-4244-0501-7/06/$20.00 ©2006 IEEE

route updates for mobile nodes. Section 3 discusses some related work. Section 4 describes our experimental setup, and presents some results in regard to EIGRP convergence as well as TCP performance. Finally, Section 5 describes conclusions and future work.

2 SIMULATION FRAMEWORK

The simulated network had to be very dynamic to effectively test EIGRP limits. This was accomplished by allowing the wireless Access Points (APs) to behave as EIGRP routers (or running EIGRP agents), and access points. Also, the end systems were allowed to retain a fixed IP address while those systems move across subnet boundaries. This way, as the mobile hosts move across network coverage, the wireless handoffs between the mobile hosts and the APs will trigger the EIGRP agents to send route advertisements to inform routers of new or revised routes to reach the mobile systems. The following sections describe our implementation efforts of EIGRP and the wireless handoff mechanism into *GTNetS*.

2.1 EIGRP

EIGRP (Albrightson, Garcia-Luna-Aceves, and Boyle 1994) is an intra-domain routing protocol that leverages the strong points of both distance-vector and link-state protocols: it converges quickly while remaining loop free at all times. This is achieved by using a system of diffused computation where every route calculation is computed in a coordinated fashion among multiple routers. EIGRP is based on the Diffusing Update Algorithm (DUAL) which is used to compute shortest paths in a distributed manner and without ever creating routing-table loops or incurring counting-to-infinity behavior. Simulation studies (Zaumen and Garcia-Luna-Aceves 1992) have shown that DUAL's average performance after a topology change (link failure, link-cost increase/decrease) is significantly better than the Distributed Bellman-Ford (DBF) algorithm used in Routing Information Protocol (RIP), and it is similar to the performance of an ideal link-state algorithm with much less CPU overhead.

EIGRP's updates are similar to a distance-vector protocol, as they are vectors of distances transmitted only to directly connected neighbors. However, the updates are partial, non-periodic, and bounded. They are partial since the updates contain only the changed routes, and not the entire routing table. They are only sent whenever a metric or topology change occurs (non-periodic), and they are sent to the affected routers only (bounded). EIGRP has shown to provide loop freedom and quick convergence in medium-scale networks(Albrightson, Garcia-Luna-Aceves, and Boyle 1994). Also, a simulation of HP backbone yielded good performance of EIGRP (Albrightson, Garcia-Luna-Aceves, and Boyle 1994). However, a true analysis and diagnosis of EIGRP protocol at grand scale has not been undertaken. This analysis is essential taking into perspective the number of deployed EIGRP-enabled Cisco routers.

We developed a scalable simulation model for EIGRP. The protocol is not ported, but rather implemented in a high quality software network simulator (*GTNetS*). Also, we have implemented a subset of EIGRP functionality as for our performance analysis, we only need link failure/restoration and link-metric change.

2.2 Wireless Handoff

As mentioned earlier, in our simulation the wireless handoffs are the events which trigger EIGRP to send routing updates. In *GTNetS* we have implemented a fairly complete subset of handoff mechanisms based on the 802.11 MAC protocol. A handoff mechanism essentially illustrates the basic steps which must be taken when a mobile station disassociates with the current access point and associates with the new one.

The wireless communications of mobile devices are vulnerable to communications interception to some degree, and thus there needs to be a control of such communications to protect the information while in transit. In our area of study, a number of security attacks could be exploited to either disrupt the functionality of the implemented protocols or to gain access to sensitive information. For example, the EIGRP updates are triggered by the wireless handoffs of the mobile nodes. Therefore a malicious user could send fake association or disassociation messages to disrupt the routing while a mobile host may or may not be already associated with an access point (the network does not converge). In addition, an adversary machine could advertise itself as an existing mobile station and associate with an access point and start receiving all packets destined to that mobile station. These packets could be very sensitive such as online banking transaction, secure access session, etc.

Due to the above reasons, we use an authentication mechanism between the APs and the mobile hosts. Once a mobile host associates with an AP, a shared key will be generated and that key will be propagated with the EIGRP updates informing the peer EIGRP routers about the new host. Thus, if a malicious user tries to send fake association or disassociation messages to disrupt the routing while a mobile host is already associated with an AP, the association/disassociation procedure will fail. This will happen because the current AP that the mobile host is associated with or any other AP has the secret key of the host being faked and knows that the message being advertised is not valid.

The wireless layer design in *GTNetS* allows for stations to be designated as Access Points (APs) or Mobile stations (MSs). The APs are connected to the wired network. Our design assumes that the APs have the role of EIGRP routers as well as access points, but this functionality can be decoupled without any effect on the proposed routing scheme. In the current scheme, we have additional local state information in the form of a last-heard timer at each MS and AP. While the AP needs to maintain one such timer for each associated MS, the MS has to maintain only one for its currently associated AP.

The handoff scheme we use is slightly different from the one proposed in the 802.11F. The mobile stations always listen to the periodic beacons sent from the Access Points (typically, every 0.1 seconds). Depending on the received signal strength (RSS) and other factors, the MS determines if the incoming beacon's transmitter is a more appropriate access point than its current association. To prevent oscillating associations we chose a threshold margin, which is the difference in signal strength that the MS must see between the current association and the incoming beacon's signal. Voluntary disassociations are initiated by sending an association request to the new AP and a disassociation message to the currently associated AP. The disassociation message is sent only after the new association has been acknowledged by an association response message by the new AP.

Both the APs and the MSs need to mutually know when each has left the operating range of the other. This mechanism is implemented by running a last-heard timer for each of the associated mobile hosts. The mobile stations send reassociation messages every 5 seconds to the access point. The receipt of a reassociation message resets the corresponding last-heard timer. A timeout of this last-heard timer means that the MS has disassociated involuntarily. On the other hand, at the MS's end, the last-heard timer is reset at the reception of a beacon from the associated AP. If the MS does not hear beacons from its currently associated AP, its last-heard timer will timeout, at which point it will assume it is no longer associated. This would make the MS try to associate with any other AP from which it hears a beacon.

The wireless MAC layer notifies the EIGRP layer of any associations and disassociations that have occurred. These notifications trigger the EIGRP diffusing computations that adjust the routing tables appropriately. For instance, when an AP receives a disassociation message from a MS (moving out of range), it sets the link metric for that MS to infinity. On the other hand, when an AP receives an association message from a MS (moving in range), it sets its link metric to a certain value (this value is the same for all MSs). We chose the value of 100, but it could be any reasonable value. Also, the secret key will be propagated with the EIGRP updates informing the peer EIGRP routers about the new MS.

3 RELATED WORK

Several simulation studies have been done to evaluate the performance of interior routing protocols for new applications/architectures. However, most of these studies tend to use OSPF models. For example, OSPF is implemented in NS-2 (McCanne and Floyd 1997), SSFNet (Cowie, Ogielski, and Nicol 2002), and GLOMOSIM (Zeng, Bagrodia, and Gerla 1998). One of the reasons could be that EIGRP is proprietary. Nevertheless, Opnet (Bertolotti and Dunand 1993) has developed an EIGRP model which could be used in such studies, and a simulation of HP backbone yielded good performance of EIGRP (Albrightson, Garcia-Luna-Aceves, and Boyle 1994). Still, one can notice that there are not enough analysis studies on EIGRP. Besides, a secondary goal of the study was to provide a mobility solution at the network layer. To our knowledge, we are the first to use route updates as a mechanism to achieve a seamless continuity of applications or sessions during mobility.

4 EXPERIMENTAL RESULTS

The test environment is a simulation of a subset of *Georgia Tech* campus which consists of 7 buildings, academic and administrative, in an area of 120 acres as shown in Figure 1(a). All 7 buildings are connected to the backbone routers with a mix of 1Gb and 100Mb links. The experimental network topology was constructed as follows. Every building network includes a wired and a wireless network. The wired network consists of 3 subnets with 30 end hosts each connected through 10Mb links as shown in Figure 1(b), while the wireless network is made of a single access point and 9 mobile stations per access point on average. The total topology has 720 wired end hosts, 72 mobile stations, and 42 routers including the 8 access points. The choice of buildings was made to include worst-case scenario EIGRP convergence. Basically, the EIGRP agents are triggered by the handoffs of the APs, and the handoffs come in two flavors: live-handoff, and dead-handoff depending whether there is an overlap in wireless coverage or not.

Kotz and Essien (2002) reported that 53 percent of the traced wireless traffic was web browsing (http/https) and the rest included data-backup, peer-to-peer file sharing, file transfer, etc. However, we believe that with today's development in wireless technology, 54g wireless cards, mobile users will tend to do most of their work through wireless media wherever possible, which leads to believe that most or all of them would have one or more long–lived active TCP connection.

We modeled two types of mobile users. First, those who start a TCP connection and remain stationary. The second are those who start a TCP connection, and then move around the campus while the connection is active. Our experiments included both types of users, as this would

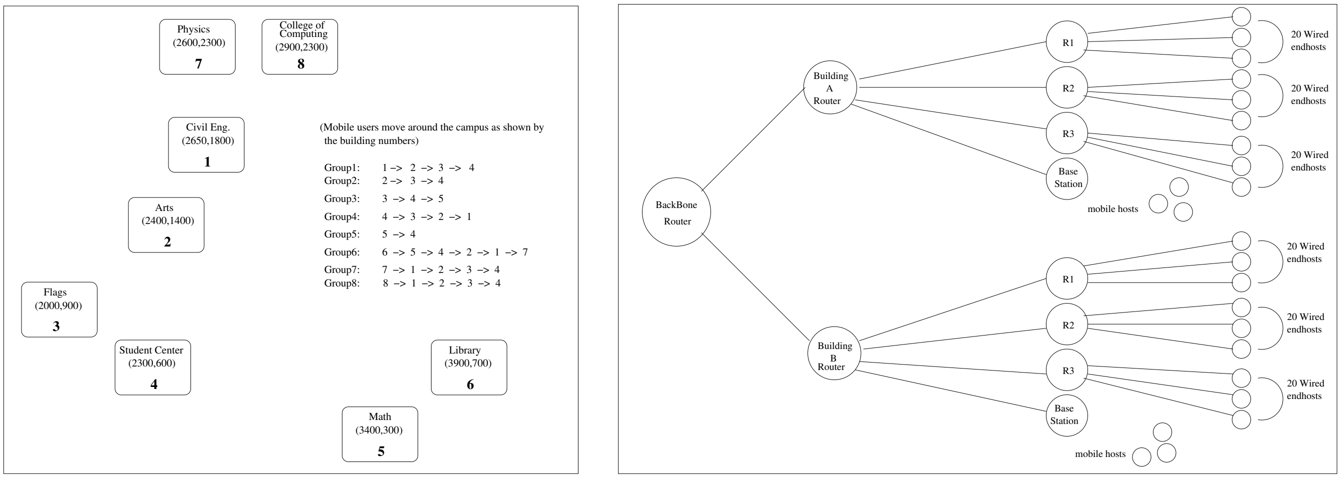

(a) Campus Building Layout

(b) Campus Building Network

Figure 1: Subset of Georgia Tech Campus

be the more realistic model and would also clearly show EIGRP convergence capability when mobile devices move across subnetwork boundaries.

All of our experiments have background traffic which includes web browsing and data-backup/file transfer traffic running on the wired end hosts. There are 700 web browsers on 200 wired end hosts that randomly connect to a group of web servers(located outside the campus), each handling a large number of simultaneous requests. The size of the individual web object requests, the size of the replies, and the time delay between the requests is modeled based on empirical measurements described by Mah (1997). There are 75 file transfers between 150 wired end hosts with a uniformly distributed size between 20 and 80 MB. The wireless users traffic is composed of 72 long-term on-campus/off-campus TCP connection uniformly distributed between 20 and 50MB each. As mentioned previously, all of the simulations were performed using our *Georgia Tech Network Simulator* (*GTNetS*), enhanced to include our detailed model of the EIGRP protocol and realistic wireless handoff models. The simulations were run for 400 simulation seconds, which resulted in a number of wireless handoff actions and routing convergence computations as reported below.

To illustrate the robustness and reliability of EIGRP, we ran four sets of experiments to collect a set of metrics. Some metrics were also chosen to demonstrate the feasibility of the mobile computing within an AS with the deployment of EIGRP. The chosen performance metrics are the following:

1. ***EIGRP Convergence Time*** is the period of time that takes the routing protocol to converge and the routing tables to reach a steady state. This metric determines the overall performance of TCP connections, since with long convergence times active TCP connections might experience substantial packet losses and several timeouts resulting in reduced performance. Figure 2 and Figure 3 show the EIGRP convergence time (log scale) throughout the simulation time for two experiments. Both experiments have the same topology shown in Figure 1(a), the only difference being the radio range for both wireless devices being 300 feet (resulting in overlapping coverage between most of the buildings) in first experiment and 200 feet (resulting in several dead zones) in the second. Initially, all EIGRP routers started randomly between 0 and 20sec, after which the mobile devices start moving according to a specific waypoint model. We see in the figures that EIGRP has a maximum convergence time of around 0.7 milliseconds for the 300ft radio range, and 9.0 milliseconds for the 200ft radio range. The convergence period is acceptable even for wireless users, as the packets destined to the MS need to be forwarded to the new AP as soon as the MS get associated with it.
It is true that the mobility model (number of users, walking patterns and speed) has a substantial effect on the EIGRP convergence time. In the worst case, handoffs would be so frequent that the protocol would never converge since during the convergence the topology has again changed. However, we believe that our experimental results are scalable to any reasonable mobility pattern and reasonable network topology size. We considered the normal walking speed(4.4ft/sec) as the general student/faculty/staff speed within the campus, and we defined specific waypoint models that best represent the mobility patterns for the campus users.

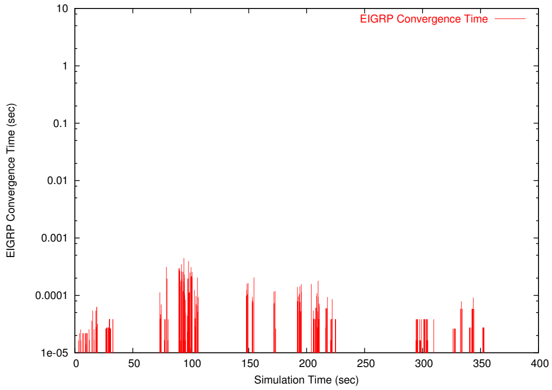

Figure 2: EIGRP Convergence Time for 300ft Radio

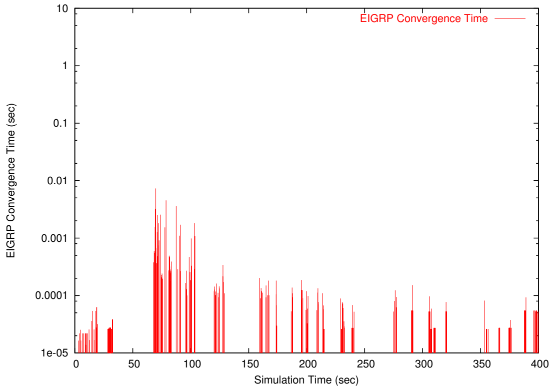

Figure 3: EIGRP Convergence Time for 200ft Radio

Since the longest measured convergence time was less than 10 milliseconds, we would have to experience more than 100 handoffs per second to overrun the convergence process with update actions. One of the reasons that our convergence times are fast is that our experiments are limited to a single AS within a small geographic region, leading to a very small propagation delay. In our wired topology, we used one microsecond for the propagation delay on all wired links. This is certainly reasonable for most moderate sized AS subnetworks.

In addition, the reason for the extra spikes in convergence time shown in Figure 3 is the number of dead zones. A dead zone occurs when a user leaves the coverage area of his existing access point association before coming into coverage range of another access point. As soon as the MS is disassociated from the AP, an EIGRP update event is triggered and results in subsequent DUAL computations. When it comes into range with a new AP, another EIGRP update event is triggered and EIGRP has to converge again.

Since most EIGRP packets are configured to have the same priority as any other packets in the network, a heavy load on the network might cause a longer convergence time for the routing protocol, due to increased queuing delay on the congested links. This is illustrated in Figure 4 and Figure 5. However, even with a large load on the network, the convergence time of EIGRP in this environment is still extremely small and acceptable for our applications.

2. ***TCP Performance*** measures the amount of data sent by an active TCP connection per unit time. The EIGRP convergence time is one of the main factors that impacts TCP performance. If the network experiences excessively long EIGRP convergence times, active TCP connections would endure heavy packet loss and numerous timeouts. Having shown with previous experiments that EIGRP in fact converged quickly due to host mobility, we expected that the TCP connections over moving wireless media with handoffs to behave as similarly as the stationary wireless media. We point out that TCP over wireless media has several other issues which are outside the scope of this paper.

The same experiments that were used to compute EIGRP convergence time resulted in the TCP performance measurement illustrated in Figure 6 and Figure 7. The figures show the TCP sequence number transmitted as a function of time. Clearly, a higher slope indicates better throughput. As mentioned earlier, we have 72 mobile devices, with each one having an active TCP connection. The flow that is shown in Figure 6 is one of the 72 wireless flows. We chose this particular one as it represents the TCP performance for a mobile station as it experiences handoffs. For this particular flow, the MS visits three networks during its mobility pattern. The handoff actions occurred at approximately 100, 200, and 300 seconds. One can notice that the first handoff was clean, with few retransmissions. However, the connection during the second handoff had no activity for short intervals. This was not a dead zone handoff; rather it was a live handoff. In cases when the mobile is at the boundary of overlapping access points, we are bound to see oscillations because of the wireless characteristics and the CSMA/CA properties. The multiple associations and disassociations that we see when the mobile station move through such a region are an artifact of the same.

The results in Figure 6 and Figure 7 are for a 300ft radio range experiment with zero and non-zero mobility. The MS shown in the bottom curve had a total of 11 handoffs. In most of our results, the stationary MS throughput was higher than the mobile one as expected. However, some of the mobile users (actively moving) had similar throughput as the stationary users as shown in Figure 7. This was

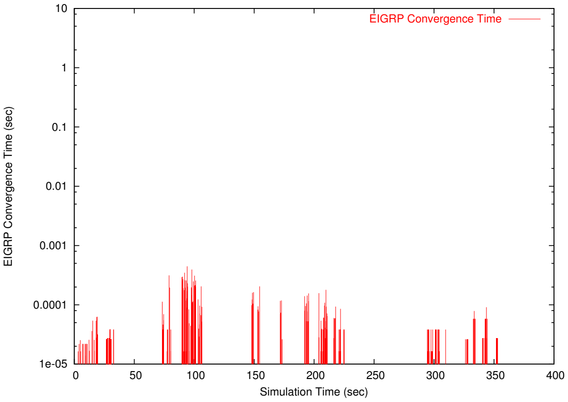

Figure 4: EIGRP Convergence Time with Data Flows

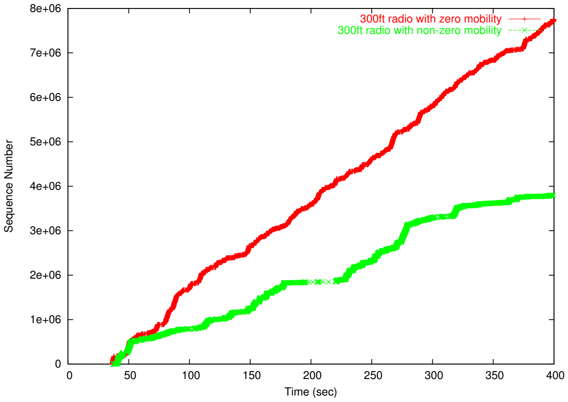

Figure 6: TCP Performance for a MS

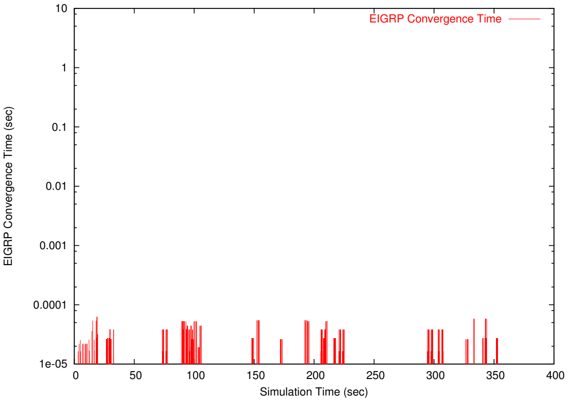

Figure 5: EIGRP Convergence Time with No Data Flows

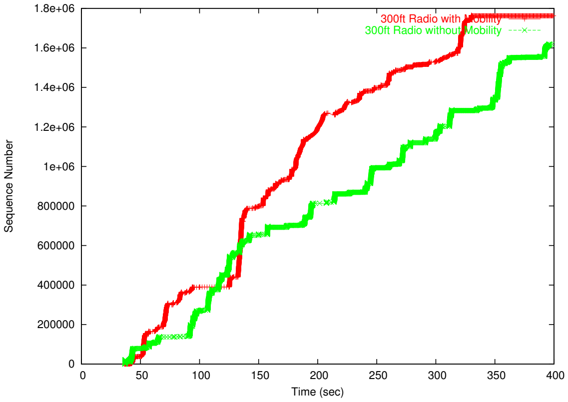

Figure 7: TCP Performance for a MS with On-campus Connection

for the on-campus connection since its short RTT (Round Trip Time) enables the wireless user with mobility to perform better. When a MS gets associated with a new AP, it resets its transmit queue. Again, the experimental results are based on the wireless users which incur a lot of dynamic network changes, thus representing worst-case scenarios for EIGRP performance.

3. ***EIGRP Overhead*** is the overhead incurred in the routing protocol due to a handoff from one access point to another. In a traditional wired network, when the EIGRP routers start they will exchange their routing table information with neighboring routers, causing many update and reply messages. After some period of time, the EIGRP protocol converges to a steady state with each EIGRP speaker having the same view of the overall network topology. As long as there is no router failure, link failure or cost metric change, there is only the low overhead of the periodic EIGRP Hello packets. However, in our experimental setup every AP is running an EIGRP protocol instance, the dynamic changes in the network due to end host mobility induce a number of EIGRP messages as the protocol recomputes the optimal routing paths. Therefore, we measured the total number of EIGRP messages to show the effect of the additional network load due to mobile handoffs and routing reconvergence. Table 1 shows the overhead of EIGRP measured when running an experiment with normal walking speed (4.4ft/sec) for mobile devices, and using a 300 foot radio range APs and MSs across the campus subset shown in Figure 1(a). When the network starts up, the EIGRP routers have to exchange the routing tables, which triggers 1,897 EIGRP update packets. After the network converges and all routing tables have reached a steady state, the MSs start their TCP connections and begin to move between the campus buildings. All MSs in our experiments follow a specific waypoint model designed to realistically model a user walking on campus. The mobility in this particular experiment resulted in an EIGRP overhead of 4,346 update packets, 4,662 query packets, and 4,662 reply packets. This may seem substantial, but recall that the EIGRP protocol uses partial updates rather than full routing table exchanges. Further, these several thousand updates were spread over a period of 400 seconds throughout our simulation execution. Since, there

are a total of 40 MS in motion and on average each moves between four buildings (according to our specific waypoint model), the 300ft radio range experiment resulted in a 284 handoffs. Some of the handoffs (a small percentage) were due to the coverage overlap, and this resulted in some oscillations.

Table 1: EIGRP Overhead (Mobility Speed 4.4 ft/sec)

	EIGRP Events During	
	Startup	Mobility
Updates	1,897	4,346
Queries	0	4,662
Replies	0	4,662

5 CONCLUSIONS AND FUTURE WORK

We have developed a scalable and detailed simulation model of EIGRP that is publicly available for researchers in the computer networks community. Also, our performance analysis of the protocol has shown its robustness and capability to adapt quickly to a very dynamic network. In addition, we have shown that the host mobility using route updates is a viable method to achieve seamless mobility and continuous connectivity for users of mobile wireless devices as they move within an AS. The EIGRP overhead incurred from mobility is minimal as all of EIGRP query and reply messages are small. Using our approach, there is no need to deploy new hosts or agents, make special configuration, or request support from any end points. We do need instances of the EIGRP protocol running on the APs, or an interface between the AP and existing EIGRP routers to inform the routing protocol of associations and disassociations.

The reliability of the experimental results depends on the accuracy of the simulation environment, and specifically the simulated network. Our results are based on a tiny network (720 wired end hosts, 72 mobile stations, and 42 routers including 8 access points.) For more realistic results, we need a campus-size simulation study (this is where the simulator scalability comes into picture). This will be our future work, where simulation will take into account the realistic deployment/placement of APs across buildings, with more than one AP in each building, and will account for true mobility models for campus users in order to provide more realism. The mobility models will include both mobility workload and user mobility model which have been characterized by a number of studies.

REFERENCES

Albrightson, B., J. J. Garcia-Luna-Aceves, and J. Boyle. 1994. EIGRP – a fast routing protocol based on distance vectors. In *Proceedings of Networld/Interop*.

Bertolotti, S., and L. Dunand. 1993. Opnet 2.4: an environment for communication network modeling and simulation. In *Proceedings of the European Simulation Symposium*.

Cowie, J., A. Ogielski, and D. Nicol. 2002. The SSFNet network simulator. `<http://www.ssfnet.org/homePage.html>`. Renesys Corporation.

Kotz, D., and K. Essien. 2002. Analysis of a campus-wide wireless network. In *Proceedings of ACM MobiCom'02*.

Mah, B. A. 1997. An empirical model of http network traffic. In *Proceedings of IEEE INFOCOMM*, 592–600.

McCanne, S., and S. Floyd. 1997. The LBNL network simulator. `<http://www.isi.edu/nsnam>`. Lawrence Berkeley Laboratory.

Moy, J. 1989. Internet RFC1131: Ospf specification. Network Working Group.

Pepelnjak, I. 2000. *EIGRP network design solutions*. Cisco Press.

Riley, G. F. 2003. The Georgia Tech network simulator. In *Proceedings of the ACM SIGCOMM Workshop on Models, Methods and Tools for Reproducible Network Research*, 5–12: ACM Press.

Zaumen, W., and J. Garcia-Luna-Aceves. 1992, December. Dynamics of link state and loop-free distance-vector routing algorithms. In *Journal of Internetworking*.

Zeng, X., R. Bagrodia, and M. Gerla. 1998, May. GloMoSim: a library for parallel simulation of large-scale wireless networks. In *Proceedings of the 12th Workshop on Parallel and Distributed Simulation*.

AUTHOR BIOGRAPHIES

TALAL M. JAAFAR is a Ph.D. candidate in the Electrical and Computer Engineering department at Georgia Tech. He also holds a M.S. in Computer Engineering from Georgia Tech. His research interests include modeling and simulation of routing protocols, large-scale simulations, and mobile computing. His e-mail address is `<jaafar@ece.gatech.edu>`.

GEORGE F. RILEY is an Assitant Professor of Computer Engineering at Georgia Tech. Dr. Riley received his Ph.D. from the Georgia Tech College of Computing in August 2001, and joined the faculty of Electrical and Computer Engineering at that time. Mr. Riley received a MSCS from Florida Tech in 1996, and a BSEE from University of Alabama in 1972. His research interests include distributed discrete event simulation techniques, large-scale computer networks, OS, and communications Front-End Processors. His e-mail address is `<riley@ece.gatech.edu>`.

DHEERAJ REDDY is a Ph.D. candidate in the Electrical and Computer Engineering department at Georgia Tech. He also holds a M. Tech. in Electrical Engineering from Indian

Institute of Technology, Madras. His research interests include modeling and simulation of wireless networks. His e-mail address is `<dheeraj@ece.gatech.edu>`.

DANA BLAIR is Technical Leader at Cisco Systems. For past 4 years, Mr. Blair has been leading the development of Optimized Edge Routing which enables Enterprises to manage multiple ISP connections for greater performance and lower cost. Prior to that, he has be contributed to Cisco's Wireless/Voice/Mobility activities and developed IOS WAN features. Prior to joining Cisco Systems in 1995, Mr. Blair worked for Hayes Microcomputer Products, IBM Network Hardware Division, and Motorola Semiconductor Products. Mr. Blair earned his Masters of Science in Computer Engineering from North Carolina State University in 1994, and Bachelor of Science in Electrical and Computer Engineering from the University of Texas in Austin, Texas in 1989. His e-mail address is `<dblair@cisco.com>`.