Improving UDP and TCP Performance in Mobile Ad Hoc Networks with INSIGNIA

Seoung-Bum Lee, Gahng-Seop Ahn, and Andrew T. Campbell, Columbia University

ABSTRACT

There is a growing need to provide better service differentiation in mobile ad hoc networks; however, this is challenging. These networks are characterized as being multihop in nature where the wireless topology that interconnects mobile hosts/routers can change rapidly in unpredictable ways or remain relatively static over long periods of time. Power and bandwidth constrained, mobile ad hoc networks typically only support best effort communications where the transport protocol's "goodput" is often lower than the maximum radio transmission rate after encountering the effects of multiple access, fading, noise, and interference. In this article we evaluate three routing protocols with INSIGNIA, an in-band signaling system that supports adaptive reservation-based services in mobile ad hoc networks. INSIGNIA represents a general-purpose approach to delivering quality of service in mobile ad hoc network supporting "operational transparency" between a number of IETF mobile ad hoc network routing protocols that include Ad Hoc On-Demand Distance Vector, Dynamic Source Routing, and the Temporally Ordered Routing Algorithm. We evaluate the performance gains delivered when using INSIGNIA with these MANET routing protocols in support of UDP and TCP traffic. The INSIGNIA ns-2 code used for the study reported in this article is available from the Web at comet.columbia.edu/insignia.

INTRODUCTION

Research and development of mobile ad hoc networks (MANETs) is proceeding in both academia and industry under military and commercial sponsorship. A number of military research projects (e.g., the Army Research Office Focused Research Initiatives, and the DARPA Global Mobile Information Systems GloMo program) are developing new MANET technologies. While a considerable amount of research is sponsored by the military, there is considerable commercial interest too. A number of companies are developing fully distributed self-configuring wireless networks that support services on demand. As a result, mobile ad hoc networking techniques are being readily applied to new fields such as sensor networks, scatter networks (i.e., interconnected personal area networks), mobile robotic networks, and deeply embedded networks. Collectively, these new technologies are promoting a world of smart spaces, and pervasive computing and communications.

Delivering services in MANETs is intrinsically linked to the performance of the routing protocol because new or alternative routes between source-destination pairs are likely to be recomputed during the lifetime of ongoing sessions. A number of efficient routing protocols have been proposed in the IETF MANET Working Group over the past several years, including Ad Hoc On-Demand Distance Vector (AODV) routing [1], Dynamic Source Routing (DSR) [2], and the Temporally Ordered Routing Algorithm (TORA) [3], among others [4]. Common features of these protocols are that they are lightweight, and provide loop-free operations and responsive routing information. The working group has focused on standardizing routing protocols suitable for supporting best effort packet delivery in IP-based networks. A number of comparisons can be found in the literature [5–8] reporting on the performance of AODV, DSR, and TORA in the context of best effort networks.

There has been little research in the area of supporting quality of service (QoS) in MANETs, however. The work that exists tends to be based on distributed scheduling algorithms that address rescheduling when the network topology changes, QoS-based medium access controllers, and fairness issues. In [9, 10], multihop multicluster packet radio network architectures are proposed. These approaches are loosely based on the virtual circuit model that requires explicit signaling for the establishment of hard state prior to communication in mobile ad hoc networks. We believe that virtual circuits lack the

0163-6804/01/$10.00 © 2001 IEEE

intrinsic flexibility needed to adapt to the dynamics found in MANETs, where the path and reservation need to dynamically respond to topology and resource changes in a timely manner. This motivates a natural separation between routing and reservation in MANETs.

If we cannot provide what is classically deemed QoS support at the networking layer (e.g., bounded delay, packet loss, and bandwidth assurances), then what level of service assurances can be delivered to applications operating in highly mobile and time-varying wireless networks? In this article we discuss the INSIGNIA signaling system [11], one proposal directed toward providing QoS support in MANETs. The INSIGNIA signaling system, which plays an important role in establishing, restoring, adapting, and removing end-to-end reservations, is a key component of a broader IP-based QoS framework for MANETs [12]. The INSIGNIA framework supports the following design features:

- Service differentiation and application adaptation
- Fast and responsive in-band signaling in support of fast reservation and restoration of services with rerouting
- Distributed resource control using soft-state resource management
- Separation between routing, signaling, and packet forwarding
- Operational transparency between multiple MANET routing protocols

INSIGNIA is responsive to changes in resource availability along communication paths and on an end-to-end basis, representing a general-purpose approach for service differentiation in MANETs; that is, INSIGNIA supports operational transparency among multiple routing protocols through the separation of routing, signaling, and packet forwarding. This is in contrast to other approaches found in the literature (e.g., [13, 14]) that call for tighter integration between resource management and routing to deliver end-to-end QoS. These approaches, however, limit operational transparency through the integration of QoS and routing.

The contribution of this article is as follows. We present an overview of the INSIGNIA signaling system, and describe our ns-2 simulation environment used for evaluation of the system. We evaluate the performance of INSIGNIA to seamlessly interoperate with AODV, DSR, and TORA, showing that the signaling system supports good operational transparency. We evaluate the performance improvement gained using INSIGNIA with the AODV, DSR, and TORA routing protocols, and present the performance improvements for UDP and TCP. Performance of the restoration algorithm relies on the speed at which routing protocols can recompute new routes between source-destination pairs when no alternative route is available after topology changes. In this case, some routing protocols outperform others in support of delivering QoS. In each case, we compare the performance of the INSIGNIA system to the baseline best effort system (i.e., AODV, DSR, and TORA without INSIGNIA) as a basis to best understand the achievable performance improvements under a wide variety of network load and node mobility conditions. We discuss our results and present some concluding remarks.

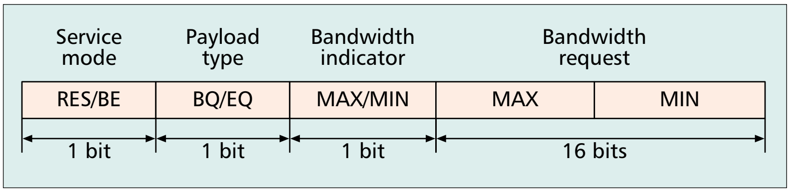

Figure 1. *The INSIGNIA IP option.*

AN INSIGNIA OVERVIEW

The INSIGNIA signaling system provides support for adaptive reservation-based services in MANETs. The signaling system supports a number of protocol commands that drive fast reservation, fast restoration, and end-to-end adaptation mechanisms. These commands are carried in-band with the data and encoded using the IP option field in datagrams. This in-band information is "snooped" as data packets traverse intermediate nodes/routers and used to maintain soft-state reservations in support of flows/microflows. As illustrated in Fig. 1, the INSIGNIA IP option supports the establishment of adaptive reservation-based services and includes service mode, payload type, bandwidth indicator, and bandwidth request fields. For a detailed description of the specification, see the INSIGNIA Internet draft [11].

FAST RESERVATION

To establish reservation-based flows between source-destination pairs, source nodes initiate fast reservations by setting the appropriate fields in the INSIGNIA IP option field before forwarding packets. A packet carrying a reservation request is characterized as having its service mode set to reservation mode (RES) and its payload set to base QoS (BQ) or enhanced QoS (EQ). Each IP packet is self-contained in that it carries all the necessary state information to establish and maintain reservations. This includes an explicit bandwidth request, as illustrated in Fig. 1. Reservation packets (i.e., data packet with the appropriate IP option set) traverse intermediate nodes, executing admission control modules, allocating resources, and establishing soft-state reservation at all intermediate nodes between source-destination pairs.

A key aspect of building QoS in MANETs is the ability of the MAC layer to deliver service quality. INSIGNIA is an end-to-end IP-based reservation mechanism designed to map down and operate over a wide variety of MAC layers. However, the stronger the assurances given by the MAC layer, the better the end-to-end performance offered to applications. In a later section we outline a modification to the IEEE 802.11 MAC distributed control function (DCF) that offers a simple set of differentiated services on which INSIGNIA is built.

A source node continues to sends packets with the reservation request bit set until the destination node completes the reservation setup phase by informing the source node of

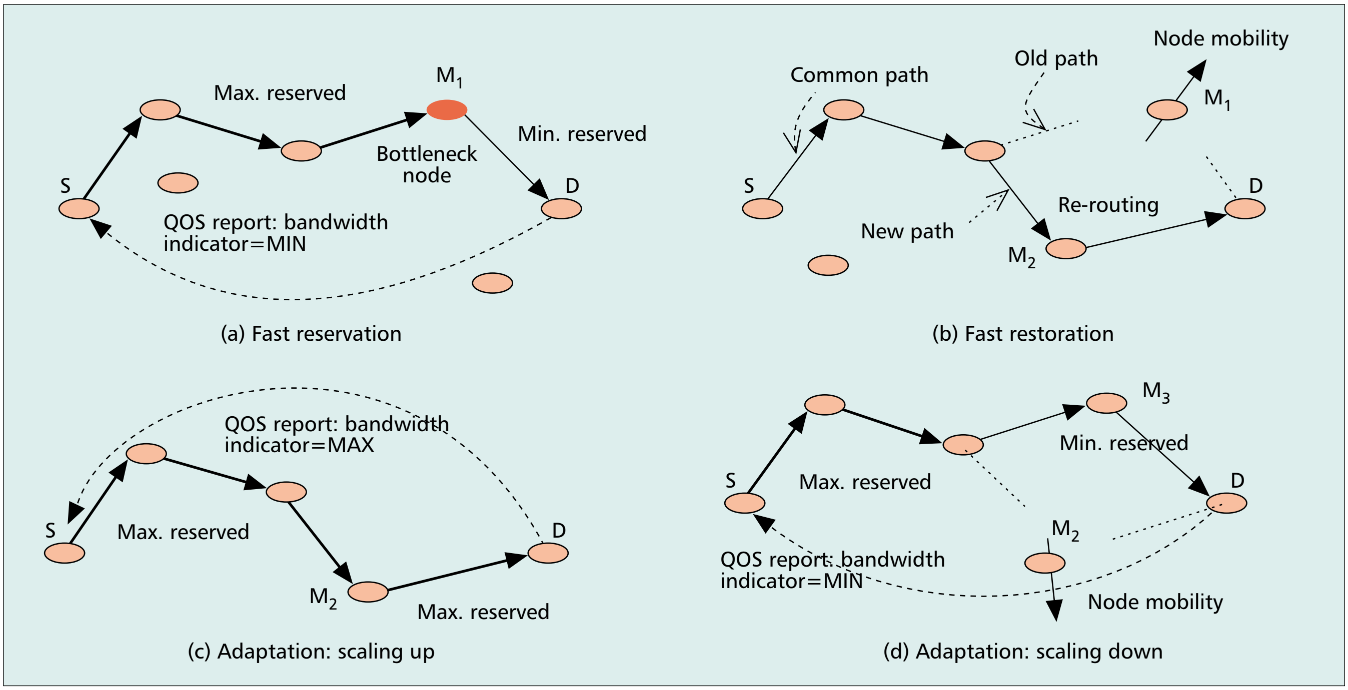

Figure 2. *Examples of INSIGNIA operations.*

the status of the reservation establishment using a QoS reporting mechanism. When a reservation packet is received at a destination node, the status of the reservation phase is determined by inspecting the service mode bit in the IP option field. The service mode bit could be set to RES for reservation or BE (best effort) for no reservation. The INSIGNIA IP option also includes a bandwidth indicator bit which can be set to MAX or MIN indicating max-reserved or min-reserved service mode, respectively. If the bandwidth indicator bit is set to MAX, it implies that all nodes between a source-destination pair have successfully allocated resources to meet the base and enhanced bandwidth requirements in support of the max-reserved service. On the other hand, if the bandwidth indication is set to MIN this indicates that only the base QoS bandwidth can be currently supported (i.e., min-reserved mode). In this case, all reservation packets with a payload of EQ that are received at the destination will have their service mode set to BE.

Figure 2a illustrates fast reservation where a source-destination pair (S, D) establishes a min-reserved flow. The destination host inspects the INSIGNIA IP option of delivered packets and determines that only a minimum reservation can be support along the current path. In this case, the base QoS packets are received with their service mode bit indicating RES, but enhanced QoS packets are delivered in best effort mode (i.e., the service mode is set to BE). The scenario shows that the bottleneck node M1 is unable to support enhanced QoS packets, and "toggles" the bandwidth indicator in the packet's IP option to MIN and sets the service-mode bit of EQ packets to BE. In this scenario the maximum reservation is provided between the source and bottleneck nodes and a minimum reservation between the bottleneck and destination nodes. We describe this as a *partial reservation*. Packets received at the destination indicate that a partial reservation has been established where only a minimum reservation service is supported on an end-to-end basis (i.e., between the source and destination nodes). The destination host informs the source node of the result of the reservation phase (i.e., minimum reservation in this case) using a QoS reporting mechanism. QoS reports traverse back toward the source node but not necessarily along the reserve path, as illustrated in Fig. 2a.

INSIGNIA is designed to operate over unidirectional and bidirectional links. However, reservations are only established on the forward link between source and destination nodes. The reception of a QoS report allows a source node to remove any partial reservation between the source and bottleneck nodes by sending EQ packets in best effort service mode; that is, by setting the EQ packet service mode bit to best effort. In this case, any resources reserved for EQ packets between the source and bottleneck nodes are automatically released by the INSIGNIA soft-state resource management mechanism, which is active at all intermediate routers.

FAST RESTORATION

Reservation-based flows are often rerouted within the lifetime of ongoing sessions due to node mobility, as illustrated in Fig. 2b. In such cases, INSIGNIA performs fast restoration. The goal of restoration is to reestablish reservations as quickly and efficiently as possible. Rerouting active flows involves the MANET routing protocol (to determine new routes), admission control, and resource reservation for nodes along the "new path." Fast restoration mechanisms also call for the removal of old reservation state at nodes along the "old path." In an ideal scenario, the restoration of a flow can be accomplished within the duration of a few consecutive packets given that an alternative route is cached. We call this type of restoration *immediate restoration*. INSIGNIA is designed to be highly responsive to node mobility in support of state restoration for rerouted flows. In essence, each IP packet is self-contained and carries sufficient state information (e.g., service mode and bandwidth request) to establish/reestablish reservations. No explicit signaling or centralized control is needed to achieve this. If no alternative route is cached, the performance of the restoration algorithm is tightly coupled to the speed at which the MANET routing protocols can discover a new path.

When a reservation-based flow is rerouted to a new node where resources are unavailable, the flow is degraded to best effort service. Subsequently, downstream nodes receiving these degraded packets do not attempt to allocate resources or refresh the reservation state associated with a flow. In this instance, the state associated with a flow automatically times out and resources are deallocated. A reservation may be restored if resources are freed up at a bottleneck node or further rerouting of flows allows the restoration process to complete. We call this type of delayed restoration "degraded restoration." If a flow remains degraded for the duration of its session, we deem it "permanently degraded."

Figure 2b illustrates a fast restoration scenario where intermediate node M_1 moves out of radio contact and a reservation-based flow is rerouted through mobile node M_2. The minimum reservation is immediately restored along the new path, while reservations along the old path are timed out and automatically removed. Note that there is no change along the *common path* as illustrated in Fig. 2b. We define the common path as any set of hops shared by the old and new paths. Resources that are freed up at nodes along the old path (e.g., at M_1) are made available to other flows. The INSIGNIA system maintains reservations through soft-state resource management. Soft-state timers are continually refreshed and reservations maintained as long as packets associated with a particular flow are periodically received at intermediate routing nodes between source-destination pairs. In contrast, if packets are not received (e.g., due to rerouting or session termination), soft-state timers expire and resources are deallocated. In the INSIGNIA system, data packets are used to maintain reservation state at intermediate nodes where the soft-state timer value is automatically coupled to the flow's data rate for optimal performance.

A major benefit of our soft-state approach is that resources allocated during the reservation phase are automatically removed in an independent and fully distributed manner when a flow's path changes due to node mobility. For example, resources at M_1 in Fig. 2b time out automatically. In this case, explicit signaling would not work because M_1 is out of radio contact with other nodes. INSIGNIA supports adaptive soft-state timer control where the reservation system "tunes" the duration of individual reservation timers to the needs of each flow in an independent fashion. Reservation-based schemes built on a soft-state resource management approach are very suitable for highly mobile environments. In [12] we report that an adaptive soft-state timer approach resolves a number of pathologies found in reservation-based MANETs, such as false restoration and resource lockup, which limit performance.

A major benefit of our soft-state approach is that resources allocated during the reservation phase are automatically removed in an independent and fully distributed manner when a flow's path changes due to node mobility.

END-TO-END ADAPTATION

The INSIGNIA system supports ongoing end-to-end adaptation that actively monitors network dynamics and adapts flows in response to observed changes based on a user-supplied adaptation policy. Flow reception quality is monitored at the destination node, and based on adaptation policy, actions are taken to adapt flows under certain observed conditions. The action taken is conditional on what is programmed into the adaptation policy by the application. For example, one adaptation policy could be to maintain the service level under degraded conditions or scale down adaptive flows to their base QoS requirements in response to degraded conditions. Another policy could be to always scale up adaptive flows whenever resources become available. The application is free to program its own adaptation policy, which is executed by INSIGNIA through interaction of the destination and source nodes.

In what follows, we describe two simple scenarios that illustrate the end-to-end adaptation process in terms of the scaling up and scaling down dynamics. The scaling up adaptation process is illustrated in Fig. 2c. Node mobility or session dynamics cause a flow routed via M_2 to be scaled up from minimum to maximum reserved service. The destination node (D) notes that the bandwidth indicator bit changes from MIN to MAX value. This indicates that the current path could support higher levels of service. This indication is really a hint from the network (and not absolute assurance) that EQ packets could be supported with reservations along the current path. In this example, resources become available at M_2, which toggles the bandwidth indicator bit of packets that traverse the node. Note that M_2 does not reserve any resources, but simply sets the bandwidth indicator bit as a hint to the destination that better QoS could be supported. It is up to the destination through interaction with the source node to use this hint to request better service. In this scenario, the destination informs the source of the resource availability via a QoS report. Based on the application's adaptation policy, the source starts to transmit EQ packets with the service mode bit set to RES. This example shows end-to-end adaptation taking place without any change in the current path between the source-destination pair. In this case, end-to-end adaptation is triggered by session-level dynamics (i.e., sessions starting, changing their bandwidth needs, or terminating) rather than mobility conditions.

The final scenario illustrates the scaling down process. In Fig. 2d a flow receiving maximum service is rerouted due to the mobility of node M_2. The new path through node M_3 has insufficient resources to support the maximum reserved

QOS is dependent on the ability of the MAC to support the end-to-end service quality semantics. While INSIGNIA is generally applicable to distributed and centrally controlled channel access schemes, we evaluate our approach within the context of existing wireless technology.

service. After restoration, the BQ packets are delivered with assurances while the EQ packets are delivered as best effort packets. The destination node (D) informs the source of this persistent degradation via a QoS report. Following this, the source scales down and starts transmitting the EQ packets in best effort mode (i.e., the service mode is set to BE). This removes the partial reservation between the source (S) and bottleneck node (M_3). Actions taken on scaling back flows is application-dependent. For example, one application may want to maintain partial reservations, hedging its bet that resources between the bottleneck (M_3) and destination (D) node will become available in the near future. Other source nodes may want to immediately remove partial reservations and forward packets in best effort mode. Some applications will not be able to tolerate best effort delivery and will scale back by dropping the EQ packets at the source node. These actions are application-specific and implemented as part of the application's adaptation policy.

INSIGNIA does not embed application-specific adaptation policy in the network (e.g., adaptation timescales, actions). Rather, it provides a simple adaptive reservation-based service model that supports service differentiation between BQ and EQ packets. Applications are free to map this service differentiation to data as they wish, monitor the network, and adapt to resource availability (by monitoring the bandwidth indicator bit) over the timescales the application considers appropriate. In essence, INSIGNIA provides a simple application programming interface (API) to the network to implement sophisticated adaptation policies at the edge (i.e., source/destination) in a scalable manner.

THE SIMULATION ENVIRONMENT

In what follows, we discuss our simulation environment used to assess the performance of UDP and TCP over INSIGNIA-enabled MANETs. The full INSIGNIA code suite and test scripts used for the evaluation of the system are freely available on the Web [15]. The simulation environment uses the ns-2 [16] simulator and its wireless extensions developed by Monarch Project [17]. In this article we use the terms *INSIGNIA system* and *best effort system* to refer to AODV, DSR, and TORA networks with and without INSIGNIA support, respectively. We present an evaluation of the best effort and INSIGNIA systems, and compare the performance of UDP and TCP traffic in both systems under diverse network load and mobility conditions.

The simulation consists of 50 mobile ad hoc nodes where each mobile node has a transmission range of 250 m and shares a 2 Mb/s radio channel with its neighboring nodes. We use a random waypoint mobility model [5] in which each mobile node selects a random destination at an arbitrary speed up to a maximum speed of 72 km/hr and pauses for a given pause time when the destination is reached. When the pause timer expires, the mobile node picks another destination and speed randomly throughout the simulation duration. The combination of pause time and velocity sets up relative degrees of mobility between mobile nodes in the simulated network. The traffic load conditions (Fig. 3) discussed in this article represent per-mobile packet generation intervals (e.g., 0.1 represents 10 packets/s per mobile host). The simulated network area has a rectangular shape of 1500 m × 300 m that minimizes the effect of network partitioning. The simulation also includes a two-ray ground reflection model and IEEE 802.11 MAC protocol.

The INSIGNIA system code [15] includes the signaling system and a number of framework mechanisms discussed in [12]. A resource monitoring mechanism allows mobile hosts to "eavesdrop" on all reserved packets within their transmission range where reserved packets represent packets associated with adaptive reservation-based flows that have passed admission control. A mobile host calibrates its estimated bandwidth availability from the bandwidth usage information snooped from reserved packets and the cached local bandwidth usage information used by a measurement-based admission control algorithm. A buffer alert mechanism is incorporated into our framework [12] to deny admission requests when a mobile node's transmission buffer and scheduler cannot accommodate new reservation requests.

As discussed earlier, QoS is dependent on the ability of the MAC to support the end-to-end service quality semantics. While INSIGNIA is generally applicable to distributed and centrally controlled channel access schemes, we evaluate our approach within the context of existing wireless technology. In [18] we describe a MAC layer based on modifications to the IEEE 802.11 distributed control function that provides simple differentiated services. The MAC ensures that not only packets sent by the mobile host itself are differentiated, but, more important, that differentiation is effective among packets sent by all other mobile hosts as well. Effective service differentiation, which is achieved in a fully distributed manner [18], is possible by appropriately adjusting the backoff times through the contention window limits. Two classes of service are supported by the MAC. The RES packets, QoS reports, and routing control messages are delivered using a high-priority service, while the BE packets are carried by a best effort MAC service. Initially, we only considered supporting RES packets using the high-priority MAC service; however, we observed that routing update and maintenance packets are often delayed and lost, causing time-consuming route updates and stale network state to persist. For this reason, we made all routing control high-priority packets. For more details on our modified MAC used throughout this study see [18].

Twenty flows are active during the simulation and are started with staggered times. Six of these flows are arbitrarily selected and monitored for a duration of 300 s in the INSIGNIA and best effort systems. The remaining flows represent cross traffic that introduces dynamic loading into the network. The traffic load ranges from 628 kb/s to 1.39 Mb/s. The network resources are partitioned a priori such that at most 800 kb/s is allocated for reservation-based flows with the remainder supporting best effort traffic. This

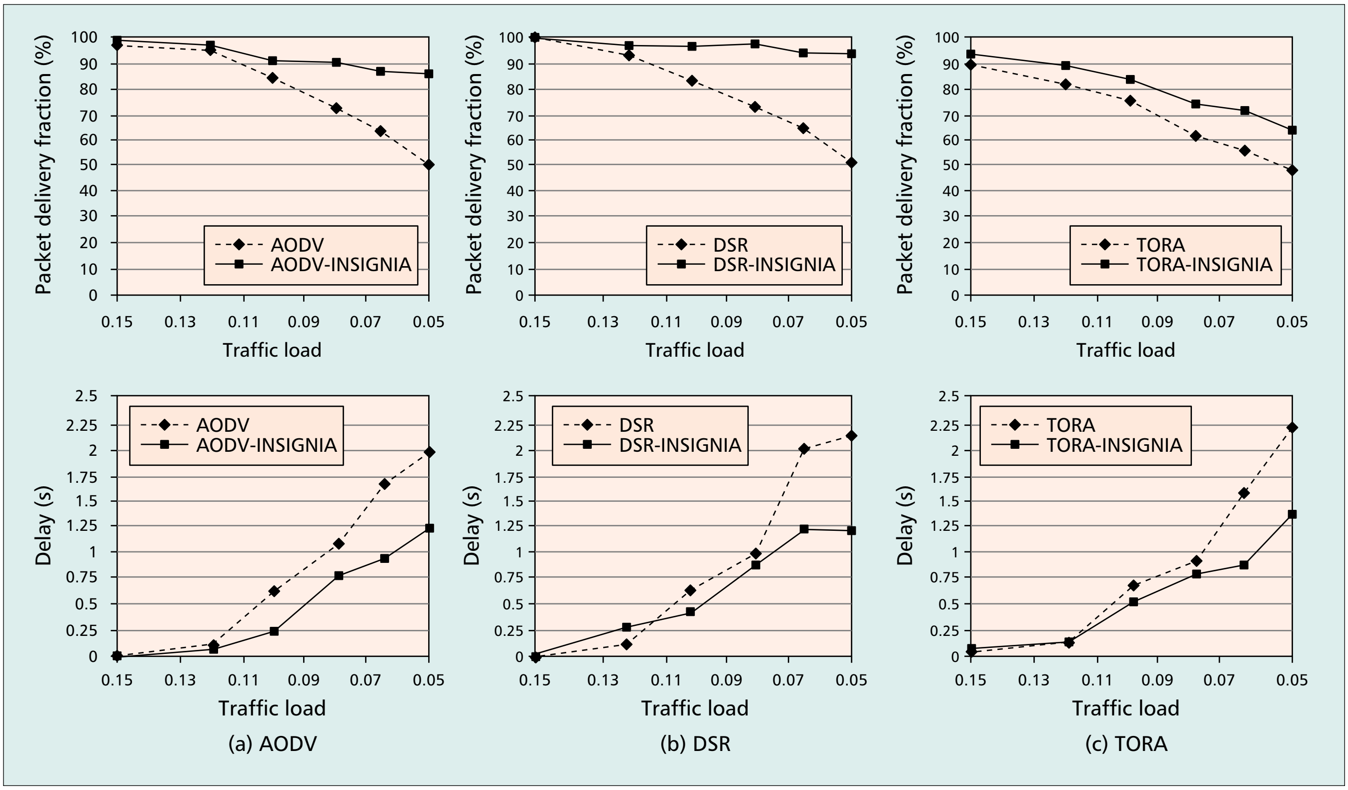

(a) AODV

(b) DSR

(c) TORA

■ **Figure 3.** *Comparison of the best effort and INSIGNIA systems under increasing network load.*

partitioning avoids starvation of best effort service packets in the presence of a large number of reservation-based flows. The various mobility conditions range from 300 s pause time, which represents no mobility, to 0 s pause time, which represents continuous mobility with a maximum speed of 20 m/s (72 km/hr). We measure a number of metrics to get an understanding of the performance of the two systems under study. These metrics include packet delivery fraction, goodput, and end-to-end delay.

In the following section we evaluate the impact of traffic load and mobility on AODV, DSR, and TORA routed networks encompassing both the reservation-based and best effort systems with particular focus on UDP and TCP performance.

UDP PERFORMANCE

Previous performance comparisons [5, 6] of AODV, DSR, and TORA in best effort networks have often favored lightly loaded networks with relatively small packet sizes. As a result, measured performance often achieves over 90 percent packet delivery fraction (i.e., the number of packets received divided by the number of packets sent). Because there is little or no congestion experienced in the simulations discussed in these comparison studies, negligible end-to-end delays are observed. These results do not hold as traffic load increases in MANETs, however. In this section we evaluate the performance of these routing protocols over a range of network conditions including heavily loaded networks with high mobility. The result is that flows often experience congestion, packet loss, and unpredictable end-to-end delays.

IMPACT OF TRAFFIC LOAD

The impact of traffic load on the performance of the best effort system in terms of packet delivery fraction and end-to-end delay is shown in Fig. 3. The x-axis represents the network traffic load in terms of UDP packet generation intervals. The traffic load is gradually increased under moderate mobility conditions (i.e., a pause time of 120 s) while the performance of the six monitored flows is observed. Identical simulations were conducted for AODV, DSR, and TORA networks to show the operational transparency of INSIGNIA working with these routing protocols and to observe the performance differences that exist among these different MANET routing protocols.

As shown in Fig. 3, the best effort system (represented by the dotted lines in the plots) achieves more than a 90 percent packet delivery fraction under lightly loaded network conditions. This is consistent with results found in the literature [5–7]. Because congestion is not evident, packets experience little delay under these conditions. However, as the traffic load increases, the packet delivery fraction decreases and the corresponding end-to-end delay increases for all of the MANET routing protocols under study. In the best effort system, the packet delivery fraction drops below 81 percent for all MANET routing protocols when the cross traffic exceeds 716 kb/s, representing a packet generation interval of 0.08 s. In addition, less than 60 percent of the packets are delivered when the cross traffic increases to 1.14 Mb/s, representing a packet generation interval of 0.05 s. Correspondingly, the end-to-end delay measurements show a substantial

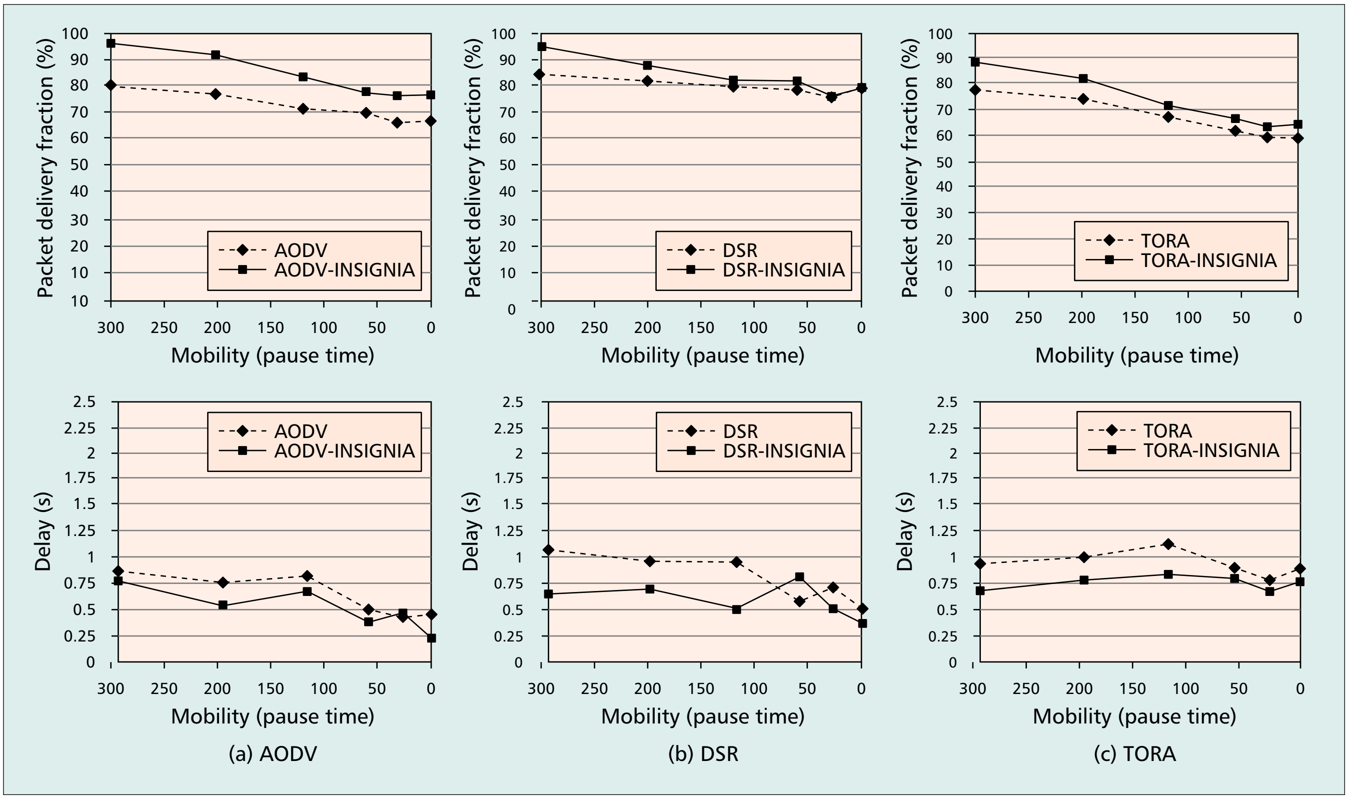

(a) AODV (b) DSR (c) TORA

■ **Figure 4.** *Comparison of best effort and INSIGNIA systems under increasing node mobility.*

increase as the traffic load increases. These results demonstrate that the delivered service quality for best effort MANET networks quickly degrades as the load of the network increases. The reservation-based INSIGNIA system provides performance improvements for UDP traffic over the best effort system, as represented by solid lines in the plots shown in Fig. 3. The performance improvements of the INSIGNIA system are shown in comparison to the best effort system for each MANET routing protocol (AODV, DSR, and TORA) under study.

As shown in Fig. 3, there is no performance gain achieved by the INSIGNIA system under lightly loaded network conditions. There is very little need for reservation in lightly loaded networks that are underutilized. However, as the traffic load increases, the INSIGNIA system outperforms the best effort system. In the case of the DSR best effort system, the packet delivery fraction drops to 91 percent when a cross traffic load of 573 kb/s (represented by a packet generation interval of 0.10 s) is introduced into the best effort system. As cross traffic load increases to 1.14 Mb/s (represented by a packet generation interval of 0.05 s), only 60 percent of the packets are delivered. In contrast, the packet delivery fractions for reservation-based flows do not drop below 88 percent for the INSIGNIA system even under heavily loaded conditions. This result is very encouraging. The improvement is due to the service differentiation supported by the INSIGNIA system where reservation-based flows are valued over best effort traffic.

The corresponding improvements in the end-to-end delay measurements are also shown in Fig. 3. We observe that under lightly loaded conditions the average end-to-end delay for the INSIGNIA system is slightly larger than that experienced by the best effort system. This is due to the additional signaling messages generated by the INSIGNIA QoS reporting mechanism. Periodic and event-based QoS reports traversing back toward the source often create additional routing information. However, the transient behavior disappears and the benefits of INSIGNIA become evident as more traffic is introduced. The average end-to-end delay under moderate to heavily loaded conditions often shows more than 80 percent improvement in the INSIGNIA system for all the MANET routing protocols, as shown in Fig. 3. We observe that AODV and DSR behave in a similar fashion as the traffic load increases in the best effort system as well as in the INSIGNIA system, while TORA slightly underperforms due to the number of signaling messages generated to create and maintain valid routes.

IMPACT OF MOBILITY

The impact of node mobility on the performance of the best effort system is shown in Fig. 4. The simulation duration is set to 300 s with 20 flows active in the network. We use the same mobility metric (i.e., pause time) defined in [5, 6] to align our simulation results. The maximum and minimum mobility conditions are represented by a pause time of 0 and 300 s, respectively. The effect of mobility is observed by gradually decreasing the pause time of mobile nodes with the traffic load fixed at 800 kb/s (i.e., 40 kb/s/flow). As shown in Fig. 4, as mobility increases, packet loss and end-to-end delay grow. One interesting observation is that the majority of packet loss is not due to loss over the wireless links. Rather, most packet loss is due to packet drops at congestion points where short-lived congestion hotspots are a result of node mobility. The IEEE 802.11 link layer retransmission scheme effectively handles packet loss over wireless links. Congestion hotspots are typically observed at intermediate mobile nodes that encounter traffic bursts after topology changes. Such conditions are very difficult to control and provision for in MANETs. This inevitably leads to degraded restoration of rerouted reservation-based flows. Increased mobility results in shorter observed congestion periods but increases the number of congestion hotspots observed in the network. In addition, faster mobility decreases the stability of routes, and consequently flows encounter fluctuations in resource availability on various paths during the lifetime of sessions. This contributes to service disruption and degradation at the destination. While many flows experience degraded service quality when mobility increases, some flows benefit from increased mobility. This is rather counterintuitive. This phenomenon is due to the effect of load balancing across the routes in the network caused by mobility. Those flows experiencing congestion under low mobility conditions improve their performance by being rerouted out of a congested portion of the network as mobility increases. This phenomenon is also observed in [7].

Figure 4 shows the impact of mobility on the best effort and INSIGNIA systems with respect to the packet delivery fraction and delay. The best effort network is limited in support of real-time applications as mobility increases. Similar trends are observed for all MANET routing protocols in the best effort network.

Figure 4 compares the performance measurements of six monitored flows in the best effort and INSIGNIA systems. The INSIGNIA network outperforms the best effort network under low to moderate mobility conditions across all routing protocols. INSIGNIA delivers at least 10 percent improvement in the packet delivery fraction for AODV and DSR under low mobility conditions and more than 7 percent for the TORA protocol. As mobility increases, the benefits of INSIGNIA over the best effort network narrow, as shown in Fig. 4. Under high mobility conditions (i.e., 72 km/hr) the INSIGNIA system provides little performance improvement over the best effort network performance. We observe that the benefit of a reservation at very high mobility is discounted by the fact that reservation holding times are very short-lived before another rerouting event occurs. In addition, the load-balancing phenomenon is observed at high mobility where flows are “spread” across the network. We also note that the signaling load increases as mobility increases in order to update/maintain routing information, decreasing available network resources.

The end-to-end delay measurements of the monitored flows in the INSIGNIA system also show improvement over the best effort system, as shown in Fig. 4. We note that there is a difference between the INSIGNIA and best effort systems in terms of the number of delivered packets. In the case of AODV, the packet delivery fraction for the INSIGNIA system is 92 percent when mobility is set at 200 s pause time in contrast to 77 percent in the best effort system. The average end-to-end delay measurement of 0.75 s in the best effort system corresponds to the 80 percent packet delivery fraction, while the average end-to-end delay measurement of 0.51 s in the INSIGNIA system corresponds to the 92 percent packet delivery fraction. The INSIGNIA system not only decreases the packet loss but also reduces the end-to-end delay.

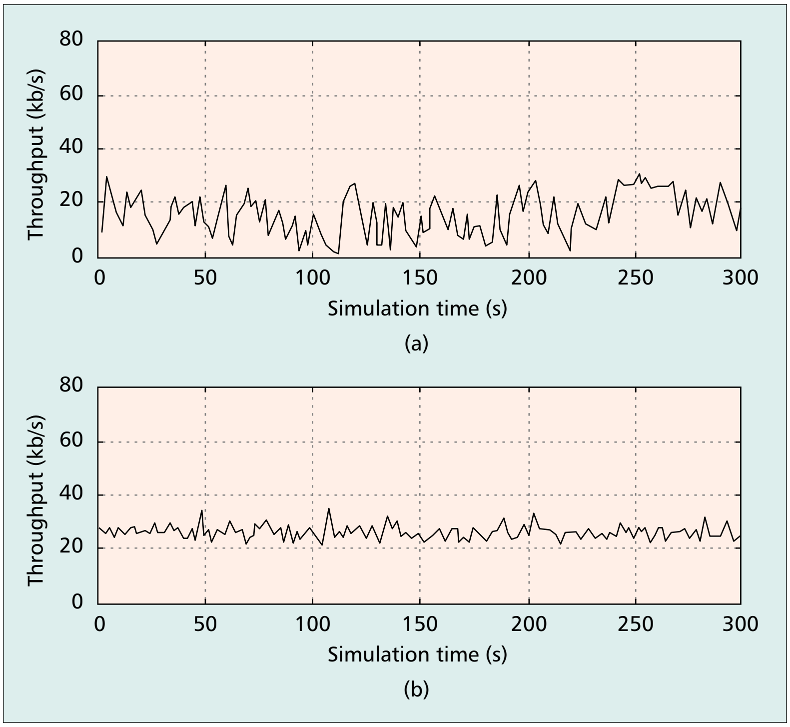

Figure 5. *a) A trace of a monitored flow in a best-effort system; b) a trace of a monitored flow in an INSIGNIA system.*

Figure 5 compares the same monitored flow under identical operating conditions in the best effort and INSIGNIA systems. The service quality measured at a destination host is shown in the figure. The throughput trace corresponds to a 30 kb/s UDP/constant bit rate (CBR) flow operating under low to moderate mobility conditions (i.e., 120 s pause time). The bandwidth requirement for the flow is defined by a minimum data rate of 22 kb/s. Figure 5a shows the throughput trace of the flow in the best effort system, and Fig. 5b shows the throughput trace of a reservation-based flow in the INSIGNIA system. The monitored flow is rerouted six times during the simulation period and traverses three wireless hops on average. Service disruption is observed on numerous occasions in the best effort trace. The throughput fluctuates throughout the trace, dropping below the minimum data rate requirement of 22 kb/s. In addition, 43 percent of the transmitted packets are lost and 65 percent of the delivered data packets exceed 800 ms end-to-end delay. In contrast, near constant rate throughput is observed for the same flow in the INSIGNIA system with 2 percent packet loss and only 9 percent of delivered packets exceeding 800 ms end-to-end delay.

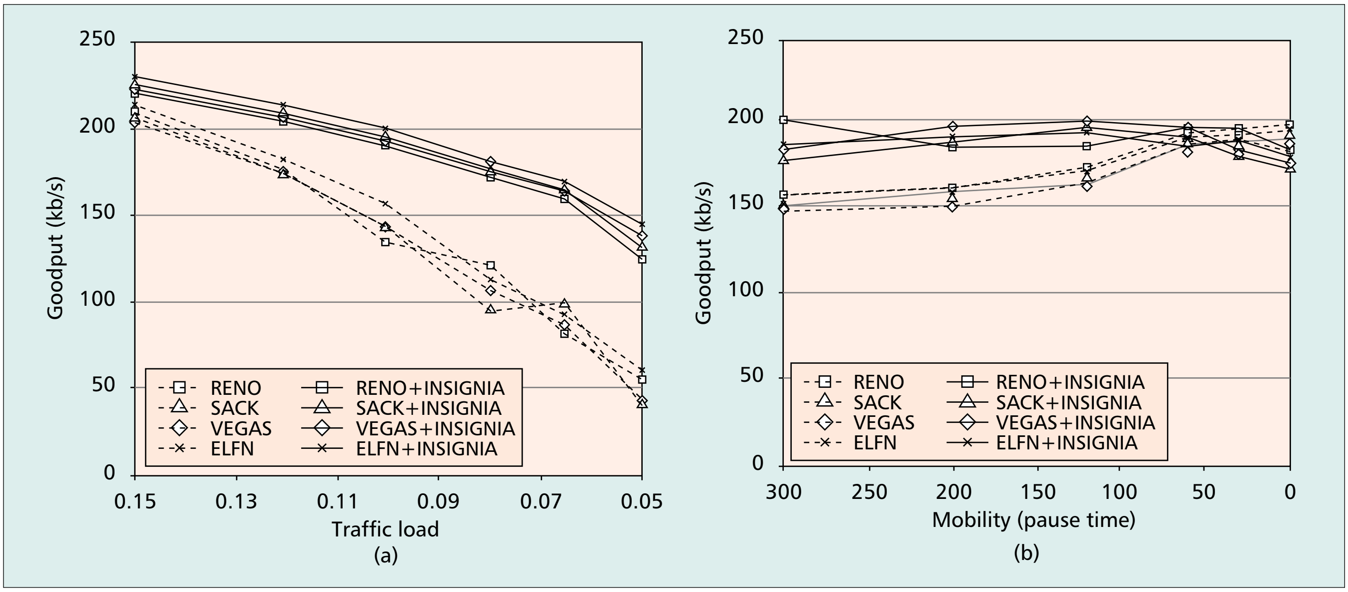

Figure 6. *a) Impact of network traffic load on TCP goodput for best-effort and INSIGNIA systems; b) impact of mobility on TCP goodput for best-effort and INSIGNIA systems.*

TCP PERFORMANCE

Most performance comparisons of MANET routing protocols have been conducted using UDP for the transport of CBR traffic. However, TCP may be the dominant transport in MANETs. The protocol behavior of TCP is quite different from UDP, embodying reliable end-to-end packet delivery and guaranteed in-order packet delivery of data to applications. Any packet loss, out-of-sequence data, or excessive delay may cause a TCP source to retransmit packets, which consequently impacts the goodput (i.e., the actual amount of data received by the destination node). Typically, TCP runs over best effort networks and configures itself to operate at the bottleneck node between source-destination pairs. In what follows, we discuss the performance of TCP for the best effort and INSIGNIA systems.

We present the performance of various TCP protocols — TCP-Reno, TCP-SACK, and TCP-Vegas —over the best effort and INSIGNIA systems. We also evaluate the explicit link failure notification (ELFN) [19], which is specifically designed to enhance TCP in MANETs.

THE IMPACT OF TRAFFIC LOAD

We observe the impact of traffic load on the six monitored TCP flows under identical network conditions to the UDP simulations discussed in the previous section. A packet size of 512 bytes and a maximum window size of 20 are used. The experiments are conducted under moderate mobility conditions (i.e., 120 s pause time). The impact of increasing traffic load on TCP-Reno, TCP-SACK, TCP-Vegas, and TCP-ELFN shows similar trends, as shown in Fig. 6a. The INSIGNIA system provides marginal improvement in goodput over the best effort system when the network load is 628 kb/s (represented by a packet generation interval of 0.15 s). However, as the network load increases the performance improvement increases, as shown in the Fig. 6a. The goodput performance of the monitored flows decreases below 70 kb/s when the traffic load increases to 1.39 Mb/s (maximum load) in the best effort system. In contrast, the goodput of the six monitored flows in the INSIGNIA system remains above 125 kb/s under maximum load. This performance improvement represents a 150 percent increase in goodput for all versions of TCPs operating at maximum load. All TCP variants operate with some differentiation, as shown in the figure.

THE IMPACT OF MOBILITY

The impact of mobility on TCP flows in terms of goodput is shown in Fig. 6b. To observe the impact of mobility on TCP goodput, we fix the traffic load at 800 kb/s and gradually increase the mobility of nodes. A traffic load of 800 kb/s is sufficient to produce congested conditions for the shared 2 Mb/s wireless channel used in our simulations. The actual bandwidth availability decreases with the number of active mobile nodes (i.e., those transmitting/forwarding packets) within each other's transmission range. For example, if two intermediate mobile nodes forwarding packets for one of the reserved flows are within each other's transmission range, the maximum available resources perceived by each mobile host are well below 1 Mb/s. The results indicate that TCP is resilient to mobility and performs well under high mobility conditions. We observe that the monitored TCP flows improve their goodputs under high mobility conditions in the best-effort system. This is a product of the load balancing phenomena discussed in an earlier section. We observed a number of different behavior characteristics across the monitored flows. Some flows encountering minor congestion experience service degradation at increased mobility, while others, experiencing congestion achieve improved goodput through rerouting brought about by node mobility.

Substantial improvements in goodput is observed at lower mobility levels where the routes are more stable and end-to-end reservation remains stable for longer periods of time. As mobility increases, the improvement of the INSIGNIA system over the best effort system narrows because the reservation holding times are short-lived before another rerouting event occurs. The INSIGNIA system not only improves TCP goodput but also shows improved service quality over all mobility conditions. At high mobility, TCP flows often decrease their window segment size to the minimum due to packet losses resulting from lack of connectivity or congestion experienced in the network. More congestion points are observed under higher mobility. Here increased mobility causes frequent topology changes often creating more bursty traffic for multiple TCP flows at a common node (e.g., a hotspot) where only limited wireless resources are available.

CONCLUSIONS

In this article we present an overview of the INSIGNIA signaling system, and evaluate the performance of AODV, DSR, and TORA operations in best effort and INSIGNIA systems. Furthermore, we discuss the performance improvements for UDP and TCP when using the INSIGNIA system. Our results confirm that INSIGNIA supports operational transparency between multiple MANET routing protocols (i.e., AODV, DSR, and TORA), and enhanced performance for UDP and TCP traffic under various node mobility and network load conditions.

The INSIGNIA system combines a number of techniques such as in-band signaling, soft-state resource management, and per-packet state management. These techniques provide a foundation for fast reservation, fast restoration, and end-to-end adaptation. INSIGNIA is responsive to the mobility of nodes, load on the network, and ability of applications to adapt. As a result, we believe that INSIGNIA is well suited to support adaptive real-time applications in mobile ad hoc networks.

In terms of future work, we are studying how best to integrate our recent work on power-aware routing (PARO) [20] in mobile ad hoc networks with the INSIGNIA system. We believe that it not only necessary to provide QoS support in mobile ad hoc networks, but we must be capable of doing so while minimizing the power consumed by mobile ad hoc devices. Finally, the INSIGNIA ns-2 code used for the study reported in this article is available on the Web (comet.columbia.edu/insignia).

ACKNOWLEDGMENTS

This work is supported in part by the Army Research Office (ARO) under ARO Award DAAD 19-99-10287. The authors would like to thank Prof. Mischa Schwartz and members of the COMET Group for input to this work.

REFERENCES

[1] C. E. Perkins and Elizabeth Royer, "Ad-Hoc On Demand Distance Vector Routing," *2nd IEEE Wksp. Mobile Comp. Sys. and Apps. (WMCSA)*, New Orleans, LA, Feb. 1999.

[2] D. B. Johnson and D. A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Network," *Mobile Computing*, Ch. 5, pp. 153–81.

[3] V. Park and S. Corson, "Temporally Ordered Routing Algorithm (TORA) Version 1 Functional Specification," draft-ietf-manet-tora-spec-00.txt, Nov. 1997, work in progress.

[4] IETF MANET Working Group, http://www.ietf.org/html.charters/manet-charter.html

[5] J. Broch *et al.*, "A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols," *4th Annual ACM/IEEE Int'l. Conf. Mobile Comp. and Net.*, Dallas, TX, Oct. 1998.

[6] V. Park and M. S. Corson, "A Performance Comparison of the Temporally-Ordered-Routing Algorithm and Ideal Link-State Routing," *IEEE Symp. Comp. and Commun. '98*, Athens, Greece, June 1998.

[7] S. Das, C. E. Perkins, and E. Royer, "Performance Comparison of Two on Demand Routing Protocols for Ad Hoc Networks," *IEEE INFOCOM 2000*, Tel Aviv, Israel, Mar. 2000.

[8] P. Johansson *et al.*, "Routing Protocols for Mobile Ad-Hoc Network — A Comparative Performance Analysis," *MOBICOM '99*, Seattle, WA, Aug. 1999.

[9] C. R. Lin and M. Gerla, "Asynchronous Multimedia Multihop Wireless Networks," *IEEE INFOCOM '97*, Apr. 1997.

[10] S. Ramanathan and M. Steenstrup, "Hierarchically-organized, Multihop Mobile Networks for Multimedia Support," *ACM/Baltzer Mobile Networks and Applications*, vol. 3, no. 1, pp. 101–19.

[11] S.-B. Lee *et al.*, "INSIGNIA," IETF Internet draft, draft-itef-manet-insignia-01.txt, Nov. 1999, work in progress.

[12] S.-B. Lee *et al.*, "INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks," *J. Parallel and Dist. Comp.*, Special issue on Wireless and Mobile Computing and Communications, vol. 60 no. 4, Apr. 2000, pp. 374–406.

[13] P. Sinha, R. Sivakumar, and V. Bharghavan, "CEDAR: A Core-Extraction Distributed Ad Hoc Routing Algorithm," *IEEE INFOCOM '99*, New York, NY, Mar. 1999.

[14] T. W. Chen, M. Gerla, and J. T. Tsai, "QoS Routing Performance in Multi-hop, Wireless Networks," *IEEE ICUPC '97*, San Diego, CA, Oct. 1997.

[15] The INSIGNIA Project of Columbia Univ., http://comet.columbia.edu/insignia/

[16] Network Simulator Notes and Documentation, UCB/LBNL, http://www-mash.cs.berkeley.edu/ns/

[17] The Monarch Project of Carnegie Mellon Univ., http://www.monarch.cs.cmu.edu/

[18] M. Barry, A. T. Campbell, and A. Veres, "Distributed Control Algorithms for Service Differentiation in Wireless Packet Networks," *IEEE INFOCOM 2001*, Anchorage, AK, Apr. 2001.

[19] G. Holland and N. Vaidya, "Analysis of TCP Performance over Mobile Ad Hoc Networks," *MOBICOM '99*, Aug. 1999.

[20] J. Gomez *et al.*, "PARO: A Power-Aware Routing Optimization Scheme for Mobile Ad hoc Networks," IETF Internet draft, draft-gomez-paro-manet-00.txt, Mar. 2001, work in progress.

BIOGRAPHIES

SEOUNG-BUM LEE (sbl@comet.columbia.edu) received his B.S. degree in electronics engineering from Yonsei University, Korea, in 1993 and his M.S. degree in electrical engineering from Columbia University, New York, in 1996. He is currently a Ph.D. candidate and member of the COMET group (http://comet.columbia.edu) and Columbia Networking Research Center (http://cnrc.columbia.edu). His areas of interests are mobile wireless networking and related QoS issues.

GAHNG-SEOP AHN (ahngang@comet.columbia.edu) received his B.S. degree in electrical engineering from Yonsei University, Korea, in 1995 and his M.S. degree in electrical engineering from Columbia University, New York, in 2000. In 1999 he joined the COMET group and Columbia Networking Research Center and is currently a Ph.D. candidate. His research interests are in QoS and mobility issues in wireless networks.

ANDREW T. CAMPBELL (campbell@comet.columbia.edu) is an associate professor in the Department of Electrical Engineering and a member of the COMET Group and Columbia Networking Research Center, Columbia University. His areas of interest encompass mobile computing and networking, open programmable networks, and quality of service research. He received his Ph.D. in computer science in 1996 and the NSF CAREER Award for his research in programmable mobile networking in 1999.

The INSIGNIA system combines a number of techniques such as in-band signaling, soft-state resource management and per-packet state management. These techniques provide a foundation for fast reservation, fast restoration, and end-to-end adaptation.