Adaptive Quality of Service for Wireless Ad hoc Networks

Seoung-Bum Lee

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

in the Graduate School of Arts and Sciences

Columbia University

2006

ABSTRACT

Adaptive Quality of Service for Wireless Ad hoc Networks

Seoung-Bum Lee

This thesis contributes toward the design of a new adaptive quality of service (QOS) paradigm for wireless ad hoc networks. We address some of the key performance problems in the broader realm of wireless ad hoc networks, including mobile ad hoc networks and emerging wireless ad hoc sensor networks.

Wireless ad hoc networks represent autonomous distributed systems that are infrastructureless, fully distributed, and multi-hop in nature. Over the last several years, wireless ad hoc networks have attracted considerable research attention in the general networking and performance community. This has been fueled by recent technological advances in the development of multifunctional and low-cost wireless communication devices. Wireless ad hoc networks have diverse applications spanning several domains, including military, commercial, medical, and home networks. The results of all this research activity the wireless ad hoc networks are starting to move from the research domain into the real world and are being gradually integrated into our daily lives. Projections indicate that this will accelerate later in the decade, to the point where some analysts predict that these types of self-organizing wireless devices will eventually become the dominant form of communications infrastructure.

To cope with the unpredictable nature of this highly dynamic environment, wireless ad hoc networks need to be able to adapt to changes in resource availability (i.e., energy, bandwidth, processing power, network density, and topology changes) and overcome any unanticipated networking problems while satisfying a wide range of application requirements. Meeting these requirements in such an environment is very challenging because the performance observed by users, devices, and routing paths selected through the network will continuously change in response to the time-varying network dynamics.

This thesis addresses some of the key issues needed to meet the requirements in support of the adaptive QOS for wireless ad hoc networks. They include (1) an adaptive QOS framework and signaling protocol for mobile ad hoc networks, (2) congestion mitigation in mobile ad hoc networks, and (3) a cost-efficient agile routing mechanism for wireless ad hoc sensor networks.

In the contribution of this thesis, we study the technical challenges for QOS support in mobile ad hoc networks and propose the INSIGNIA QOS framework that is designed to support the adaptive service paradigm. The key component of the QOS framework is the *INSIGNIA signaling system*, an in-band signaling system specifically designed to address the adaptive QOS related challenges in mobile ad hoc networks. The INSIGNIA signaling system is recognized as one of the first QOS signaling protocols in mobile ad hoc networks. We also present a detailed performance evaluation of the IEEE 802.11 based INSIGNIA signaling system with a number of well-known MANET routing protocols. The INSIGNIA system shows operational transparency to a number of MANET routing protocols and offers significant performance gains for various TCP and UDP flows.

Next, we investigate the MANET-specific congestion conditions called *hotspots*. A hotspot is defined as a node (or a group of nodes) experiencing flash congestion conditions or a period of excessive contention conditions in wireless ad hoc networks. Hotspots can exist even in lightly loaded ad hoc networks and can severely degrade the network performance. The existence of a hotspot is largely due to mobility in mobile ad hoc networks and related traffic patterns where the node mobility continuously changes the network topology and causes the on-going traffic to reroute. This effect varies the network loading conditions and produces transient congestion. These hotspots cause packet loss, increase in end-to-end delay, and even trigger route maintenance as they are often misinterpreted as routing failures. As a solution to this problem, we propose a *Hotspot Mitigation Protocol (HMP)* that works with best effort routing protocols. The HMP suppresses and disperses new/rerouted flows from hotspot regions to mitigate congestion conditions. HMP also provides a traffic throttling scheme that rate controls best effort TCP flows to relieve congestion.

In the final contribution of the thesis, we shift our research focus to wireless ad hoc sensor networks, a new emerging frontier in wireless ad hoc networks. Based on the observation that current routing algorithms for sensor networks yield poor information delivery (i.e., poor fidelity as measured at the Internet gateway to the sensor network – typically called a sink), we investigate the problem and the solution space using the TinyOS embed operating system in an experimental testbed of Mica2 mote sensors. We show that the poor fidelity is largely due to the unresponsive nature of the route selection convention commonly in use in sensor networks. To resolve this problem, we propose an agile, cost effective, and high-fidelity yielding hop-by-hop routing protocol called *Solicitation-based Forwarding* (*SOFA)*. SOFA achieves fast path convergence at network deployment time and acquires an alternative path quickly with minimal signaling overhead when faced with path changing conditions due to network dynamics. Path maintenance in SOFA is minimal and when a new sensor is added to the network, it is integrated quickly and seamlessly. SOFA shows significant reduction in energy consumption where the energy savings in SOFA network are primarily due to decrease in the signaling overhead. The on-demand nature makes SOFA cost effective; its agile self-adapting nature makes it resilient to network vagaries; and its use of timely solicitation-based handshakes make its forwarding decisions effective in data delivery.

Table of Contents

1 Introduction

2 INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks

List of Figures

Acknowledgements

I am indebted to many individuals for their care and support given to me during my doctoral studies. First of all, I would like to express my deepest gratitude to Professor Andrew T. Campbell. As my advisor, he has provided me constant encouragement, insightful comments, and invaluable suggestions, which benefited not only the completion of this thesis but also molded a roadmap for my future research. I really enjoyed working with him. His hard working attitude and high expectation towards research have inspired me to mature into a better researcher. He is truly my role model.

Next, I would like to express my sincere thanks to Professor Mischa Schwartz for many insightful comments on my work and for also serving on defense committee. I would also like to thank Professor Henning Schulzrinne, Professor Shivendra Panwar, Professor Xiaodong Wang and Professor Angelos Keromytis for kindly taking time out from their busy schedules to serve on my dissertation committee.

I want to also thank Dr. Bo Ryu for giving me the opportunity to work with him in 2001 at HRL. He has shown me what the industry and research laboratories are all about and inspired me to investigate practical research. He has become a boss and a friend.

I would next express my gratitude to all my friends who made my stay at Columbia enjoyable. Many members of the COMET and CNRC group have contributed to my research and have made my time at Columbia University a truly memorable one.

Special thanks to Gahng-Seop Ahn and Sanghyo Kim. We have spent many years together simulating/implementing what seemingly never working simulator/testbed. I also like to thank Shane B. Eisenman for being such a wonderful colleague in my latter phase of my Ph.D. I will always remember and treasure the effort and input he has given me. I also thank Jiyoung Cho and Kyung Joon Kwak for helping me implementing the Hotspot Mitigation Protocol and Solicitation-based Forwarding Algorithm.

My special thanks should go to my parents since they always have loved me, believed in me, and encouraged me in my study. I am also grateful to my parents-in-law, my brother, my sister, and brother-in-law for their encouragement and prayer.

My deepest and final acknowledgment goes to my dear wife Heeyoung Kim for her dedicated sacrifice, support, understanding, and encouragement. She has endured my erratic and hectic work schedule during my Ph.D. and raised our beautiful daughter Sujin Lee. This dissertation could not be completed without her presence beside me.

Chapter 1

Introduction

1.1 Overview

The innovation in mobile computing technology and the proliferation of communication devices (e.g., cell phones, laptops, personal digital assistants, or wearable computers) are revolutionizing our way of sharing information. We are at the verge of entering the ubiquitous communication era in which a user utilizes numerous devices through which he can access all the required information whenever and wherever needed. The nature of ubiquitous communication advocates wireless networks as the most appropriate solution and as a consequence, the wireless networking realm has undergone exponential growth in the past decade.

The earliest wireless networks, called "packet radio" networks, were sponsored by Defense Advanced Research Projects Agency (DARPA) in the early 1970s. It is interesting to note that these early packet radio systems predate the Internet, and indeed were part of the motivation of the Internet Protocol. In the 1980s, DARPA continued the experiment through the Survivable Radio Network (SURAN) [1] project and endeavored to develop more sophisticated packet radio protocols that could survive electronic attacks. Another wave of academic activity started in the 1990s with the advent of inexpensive IEEE 802.11 [50] radio cards for personal computers.

Mobile ad hoc networks (MANETs) are complex distributed systems comprising wireless mobile nodes that can self-organize dynamically into arbitrary and temporary, ad-hoc network topologies. Since the mobile devices are free to move randomly, the network's wireless topology may change rapidly and unpredictably. The communication in a mobile ad hoc network can occur directly between mobile nodes or through intermediate nodes acting as routers. Minimal configuration and quick deployment make mobile ad hoc networks suitable for emergency situations like natural or human-induced disasters, military conflicts, emergency medical situations, etc, where the wired network is not available and mobile ad hoc networks can be the only viable means for communications and information access. Also, mobile ad hoc networks are now beginning to play an important role in the civilian realm (e.g., campus recreation, conferences, electronic classrooms, and in the form of various mesh networks). The introduction of technologies such as the Bluetooth, HyperLAN, GPRS (General Packet Radio System [93]), IEEE 802.11 [50], IEEE 802.15 [53], and IEEE 802.16 [54] are also fostering MANET deployments outside the military domain.

More recently, advances in micro-electro-mechanical systems (MEMS) technology have enabled autonomous wireless sensor networking of low-cost, low-power, multifunctional sensor devices. Each sensor device is capable of short distance wireless communication and has some level of intelligence to process signals. This attribute makes a sensor network highly adaptable so that it can be deployed in many environments. A wireless sensor network can be viewed as the last-mile of wireless networks, in which sensors are used to gather the desired information. However, unlike MANETs, which are predominantly used for peer-to-peer communications, the information gathered in a sensor network is typically sent directly to sink gateways (i.e., data collection entities). The decrease in the size and cost of sensors represents a new network paradigm, where a large set of "disposable" unattended sensors is used to gather, process, and deliver information. Due to limited capabilities of sensor devices, there is an extreme emphasis on energy and bandwidth conservation, which in turn, motivates the innovations in current sensor networking technologies.

Wireless sensor networks, together with mobile ad hoc networks, are part of the broader wireless ad hoc networks. Recently, these wireless ad hoc networks have come into prominence because they hold the potential to revolutionize many segments of our life, from daily communications, to military and environmental applications. However, numerous technical barriers still remain and they must be resolved before we can realize the full potential of wireless ad hoc networks. We argue that components in wireless ad hoc networks should be made adaptive and responsive to changes in network topology, node connectivity, and end-to-end quality of service conditions.

This thesis focuses on three important adaptive service enablers for wireless ad hoc networks. First, we present our view of the adaptive QOS framework and a QOS signaling protocol for MANETs. Then, we investigate the mechanisms to relieve the congestion problems in MANETs within the context of realizing the adaptive service. As the third enabler for adaptive service, we propose a cost-effective and high-fidelity yielding routing algorithm for sensor networks that provides enhanced information delivery.

1.1.1 Quality of Service in Mobile Ad hoc Networks

A mobile ad hoc network can be seen as an autonomous system or a multi-hop wireless extension to the Internet. As an autonomous system, MANET should provide its own routing protocols and network management mechanisms. As a multi-hop wireless extension, it should provide a flexible and seamless communication among the users or access to the Internet. Recently, due to increasing popularity of multimedia applications and pending commercial deployment of MANETs, the quality of service (QOS) support in MANETs has become an important requirement. However, the QOS support in a MANET is unlike that of the wireline network or the cellular network because wireless bandwidth is shared among neighboring nodes and the network topology continuously changes with node mobility. This condition requires extensive collaboration between the nodes, both to establish the route and to secure the resources necessary to provide the QOS.

According to RFC2386 [2], QOS is defined as a set of service requirements to be met by the network while transporting a packet stream from source to destination. Intrinsic to the notion of QOS is an agreement or a guarantee by the network to provide a set of measurable pre-specified service attributes to the user in terms of delay, jitter, available bandwidth, packet loss, and so on. As in the Internet, mobile ad hoc networks are designed to support the best-effort service with no guarantees of associated QOS. Therefore, when a packet is lost in a mobile ad hoc network, the sender simply retransmits the lost packet. This is an efficient method for applications requiring no QOS, but simple end-to-end retransmission is inadequate for real-time applications that are sensitive to packet loss, delay, bandwidth availability, etc.

Although a lot of work has been done in supporting QOS in the Internet, they are not readily applicable to MANET due to their resource constraints and frequent topology changes. For example, current QOS routing algorithms for the Internet require accurate link state and topology information, but the time-varying capacity of wireless links and mobility make it almost impossible to provide accurate global information in MANETs. Knowing these limitations, researchers are attempting to provide new QOS components tailored to MANETs. This research effort includes QOS routing, QOS signaling schemes (e.g., resource reservation), QOS-based MACs, and so on.

The QOS routing is different from the resource reservation (i.e., QOS signaling) and they have two distinct responsibilities that can be either coupled or decoupled in QOS architectures. The QOS routing protocol is used to find a path that meets the QOS needs, but it is the QOS signaling that reserves, maintains, and releases resources in the network. The QOS signaling will work better if it coordinates with QOS routing but most QOS routing algorithms are too complicated or too expensive (i.e., substantial overhead) to be implemented in MANET. The QOS signaling still works even without support of a QOS routing but the resource reservation may fail because the selected path may not have enough resources.

As of now, the Resource ReSerVation Protocol (RSVP) [3] and Session Initiation Protocol (SIP) [92] are the widely accepted standard signaling protocols for the Internet. However, they are not directly applicable to MANETs because the signaling overhead is too high when network topology changes. These signaling control messages will contend with data packets for the channel and cost a large amount of bandwidth. In addition, RSVP and SIP are not adaptive enough for MANETs because they have no mechanism to rapidly respond to the topology change. In particular, when the network topology changes, the signaling entity that has to manage resource reservations in the network often fails to de-allocated resources on the old path due to lack of connectivity to the targeted nodes .

The QOS MAC protocol is an essential component in QOS support in MANETs. All upper-layer QOS components (i.e., QOS routing and QOS signaling) are dependent on the QOS MAC and the ability to provide QOS is dependent on how well the resources are managed at the MAC layer. Although many MAC protocols (e.g., MACA [4], MACAW [5], FAMA [6], MACA-BI [7]) have been proposed for wireless networks, they are primarily designed to solve medium contention, hidden/exposed terminal problems but do not incorporate the notion of QOS. Recently, the Group Allocation Multiple Access with Packet-Sensing (GAMA-PS) protocol [8] and the Black-Burst contention mechanism [9] have been proposed to support QOS guarantees to real-time traffic in a distributed wireless environment. However, their QOS support is valid only in a wireless LAN environment where every host can sense each other's transmission without any hidden terminals. In fact, all aforementioned MAC protocols show some level of inadequacy for QOS support or multi-hop wireless networking. Consequently, the IEEE 802.11 [50] is the de facto standard MAC for MANETs. Various research studies have proven that the IEEE 802.11 is capable of supporting multi-hop wireless networking, effectively eliminates the hidden terminals, and provides the collision avoidance feature through its distributed control function (DCF). However, the IEEE 802.11 DCF only supports best effort service. To incorporate the notion of QOS, several researchers have proposed some modifications [65] [74] to the IEEE 802.11 DCF to support differentiated service.

Note that the signaling protocol is the control center that coordinates the behaviors of routing, MAC, and other components (e.g., admission control, scheduling). Hence, better QOS can be provided if the signaling component coordinates with other QOS modules. However, since realization of QOS components such as QOS routing and QOS MAC are often prohibitively complex and impractical in MANET environment, we need to consider implementing generic QOS measures that are not reliant on a QOS routing or a specific MAC.

1.1.2 Congestion in Mobile Ad hoc Networks

Traditionally, congestion occurs when the total volume of traffic offered to the network or part of the network exceeds the resource availability. Congestion typically manifests itself in excessive end-to-end delay and packet drops due to buffer overflow. There are a variety of conditions that can contribute to congestion and they include but are not limited to traffic volume, the underlying network architecture, and the specification of devices in the network (e.g., buffer space, transmission rate, processing power, etc).

As in the Internet, a mobile ad hoc network is also afflicted by diverse degrees of congestion. However, the cause and characteristics of congestion conditions in MANET are somewhat different from that of the Internet. This was discovered while evaluating the performance of the QOS signaling protocol discussed in this dissertation. This observation led us to study the simulation results and testbed experiments for the identification and the solution for the congestion conditions in MANET.

It was observed that the many congestion conditions in MANET are not necessarily due to the presence of excessive workloads in the network. In fact, we can observe congestions under all loading conditions, even in the lightly loaded networking condition. After a careful study of this intriguing phenomenon, it was found that the route selection convention widely implemented in MANET routing protocols is one of the key reasons for these peculiar congestion conditions.

Being a mobile network, the network topology of a MANET may change and cause a flow to reroute multiple times during the lifetime of an on-going session. The route discovery or rerouting procedure of many on-demand MANET routing protocols allows intermediate nodes to reply to route requests leading to a small number of routes becoming overused throughout the network. The mechanism to reduce the impact of flooding caused by route request packets inadvertently fosters a small number of routes to be overused, creating a unique congestion condition in MANET. In fact, we observe patches of heavily congested areas in MANETs that entail packet loss, delay spikes, and unbalanced resource consumption.

While some researchers have broadly discussed congestion issues [10] in mobile ad hoc networks, there is no comprehensive approach to this problem. This led us to investigate a generic solution that can be applied to all existing MANET protocols.

1.1.3 Challenges in Sensor Networks

As with many technologies, defense applications have been a key driver for research and development in sensor networks. The history of sensor networks started in the Cold War era with the Sound Surveillance System (SOSUS), an acoustic sensor system on the ocean bottom that was deployed at strategic locations to detect and track quiet Soviet submarines. Over the years, more sophisticated acoustic networks have been developed for submarine surveillance and SOSUS is now used by the National Oceanographic and Atmospheric Administration (NOAA) for monitoring events in the ocean, e.g., seismic and animal activity [11].

Recently, there have been great advances in MEMS technology, wireless communications, and digital electronics that have enabled the development of small sized, low-cost, low-power sensor devices that are capable of communicating short distances. These tiny multifunctional sensor devices leverage the idea of sensor networks based on the collaborative effort of a large number of nodes. Due to its size and cost, a sensor network can be deployed anywhere and it is suitable for many applications, ranging from military applications to industrial applications (e.g., for sensing, monitoring, and tracking).

However, sensor networks in general pose considerable technical challenges. Due to the potentially harsh, unpredictable dynamic environment, along with energy and bandwidth constraints, a sensor network is expected to be confronted with numerous networking problems. Sensor networks must deal with limited resources that are often dynamically changing and should operate autonomously, changing its configuration as required. The network also needs to overcome technical barriers caused by unreliable communication links that are easily affected by interferences and provide the required reliability. There have been various research efforts conducted in sensor networks focusing on several key issues; they include but are not limited to the following:

- Data dissemination and routing research [12] [13] [14] [15] for data propagation and routing in wireless sensor networks.
- Efficient sleep/duty cycle schemes [16] [17] to provide energy saving while maintaining network connectivity.
- Collaborative signal and information processing researches [18] [19] for reliable event detection and distributed information fusion,
- Energy-efficient MAC schemes [20] [21] [22] for low-power, energy conserving, or energy-efficient communications,
- Distributed time synchronization schemes [23] [24] and lightweight geographic localization mechanisms [25] [26],

- Distributed tasking and querying techniques [27] [28] for query and task compilation/placement, data organization, and caching.

As of today, overcoming these technical challenges and deploying a large-scale multifunctional sensor network is still a daunting task. However, research is underway to solve these challenges, and technical advances in sensor devices are continuously inching us towards the vision of realizing the pervasive network.

1.2 Problem Statement

This thesis investigates three fundamental issues found in wireless ad hoc networks (i.e., mobile ad hoc network and wireless sensor networks). First, we address the problem of quality of service in mobile ad hoc networks. We argue that the QOS requirements in MANET are quite different from that of the classical approach and existing QOS works predominantly designed for the traditional networks are unfit for MANETs. These QOS provisions are derived from wireline networks where the control and signaling rely on a circuit model that requires explicit connection management and the establishment of hard-state in the network prior to communication. However, out-of-band signaling needs to maintain source route information and respond to topology changes by directly signaling intermediate routers on an old path to allocate/free radio resources. In many case, this is impossible to do if the affected router is out of radio contact from the signaling entity. By the same token, the hard-state approach lacks flexibility to adapt to the dynamics found in mobile ad hoc networks. Based on this analysis we propose a new QOS architecture that can provide fast reservation, responsive restoration and seamless adaptation to mobile ad hoc network dynamics.

In Chapter 4, we address the issue of congestion conditions in wireless ad hoc networks called ‘hotspots’. We observed the hotspots in the process of studying the QOS issues presented in Chapter 2 and Chapter 3. We have identified that these conditions are specific to ad hoc networks and typically observed using on-demand routing protocols. After thorough investigation and detailed analysis of extensive simulation data, it is determined that these conditions can be prevented, identified, and mitigated (when created) through fairly simple manipulations at the routing and MAC layer. We argue that Hotspot Mitigation Protocol (HMP) is one of the generic approaches to address hotspots in MANET.

In the fifth chapter, we shift our research efforts to the realm of wireless sensor networks. We investigate the reasons for poor fidelity [29] in de facto standard routing mechanisms [15] implemented in sensor networks. Our detailed study has found that the commonly practiced routing decision (i.e., based on link estimation [15]) has very slow path convergence so that the path creation takes a long time (i.e., often tens of minutes) and when faced with network dynamics (i.e., node join, node death, temporary interference) the encountered conditions (e.g., loss of connectivity) are often undetected for long time. As a consequence, it requires an even longer time to resolve the networking impairments and this condition can result in long disruption of information delivery. We argue that components in sensor networks should be made more agile to facilitate faster adaptation to network dynamics and topology changes.

1.3 Technical Barriers

Wireless ad hoc networks have distinct system characteristics and constraints that are significantly different from traditional networks. In what follows, we discuss the important technical barriers in realizing QOS in MANETs, preventing and mitigating congestions in MANETs, and attaining high-fidelity senor networks.

1.3.1 QOS Support for Mobile Ad hoc Networks

Although substantial work has been done for quality of service in the Internet, none of the existing proposals can be readily applied to mobile ad hoc networks due to limitations and constraints intrinsic to MANETs. Supporting QOS requires the link state information such as delay, bandwidth, cost, loss rate, and error rate to be available and manageable. However, satisfying these requirements is very challenging in MANET because the quality of a wireless link can abruptly change with the dynamic of surrounding circumstances.

The traditional QOS approaches are loosely based on the virtual circuit model that requires explicit connection management and the establishment of hard-state in the network prior to communication. The virtual circuit model also assumes the route and the reservation between source-destination pairs remain fixed for the duration of a session. However, the virtual circuit lacks the intrinsic flexibility needed to adapt to the dynamics found in mobile ad hoc networks where the path and reservation need to dynamically respond to topology and resource changes in a timely manner.

Similarly, out-of-band signaling is not suitable for supporting QOS in mobile ad hoc networks. The out-of-band signaling systems are incapable of responding to fast time-scale dynamics because out-of-band signaling systems require maintenance of source route information and respond to topology changes by directly signaling affected mobiles to allocate/free resources. In some cases, this is impossible to do due to lack of connectivity between the affected router and the signaling entity that attempts to de-allocate resources over the old path.

One of the most challenging aspects of QOS support in mobile ad hoc networks is in the maintenance of service level. QOS in mobile ad hoc networks is intrinsically linked to the performance of the routing protocol because a flow between a source-destination pair is likely to be rerouted during the lifetime of on-going session. Since a rerouted flow is likely to encounter different resource availability on the new path, the QOS agreement from the old path may not be sustained any longer. The traditional assumption that the route and the reservation remain fixed for the duration of a session is no longer valid in mobile networks. Therefore, the service paradigm for mobile ad hoc networks should be adaptive in nature. Clearly, there is a pressing need for a new QOS framework for MANET that can support adaptive QOS as to dynamically respond to topology changes in a timely manner.

1.3.2 Mitigating Hotspots in Mobile Ad hoc Networks

Hotspots represent transient but highly congested regions in wireless ad hoc networks. We define hotspots as nodes that experience flash congestion or excessive contention conditions over longer time-scales. Hotspots are commonly observed in a mobile ad hoc network that implements MANET routing protocols and contention-based MACs (e.g., IEEE 802.11). Interestingly, the development of a hotspot is not necessarily related to the total traffic volume in the network but closely associated with locality of data flows that are channeled by the routing protocol. In other words, the routing protocol dictates the existence and nature of hotspots in a mobile ad hoc network. Note that hotspots are often transient in MANET because network topology continuously changes and varies the traffic loading conditions, causing hotspots to migrate. Hotspots are known to increase end-end-delay and packet loss that generally degrade the network performance. Moreover, hotspots are also the main culprit for burst of delay spikes that are often misinterpreted as loss of connectivity. The incorrect conclusion of connectivity loss often triggers mobile hosts to generate waves of routing maintenance messages that further exacerbate the already taxing conditions. Therefore, hotspot conditions should be prevented if possible and any existing hotspot conditions should be mitigated responsively.

Prevention of a hotspot involves dispersion of traffic loads from the congestion-prone area (e.g., where traffic is building up fast). Once a hotspot is created, the event has to be immediately and accurately detected followed by execution of hotspot mitigation procedures. Therefore, hotspot identification is an integral part of the hotspot mitigation procedure where the accuracy of the hotspot identification has a substantial impact on the outcome of hotspot mitigation endeavor.

There are several indicators of a hotspot in MANETs. A hotspot typically entails consecutive packet loss, excessive increase in the medium access time (MAC delay), and buffer overflow. However, identification of a hotspot based on a single indicator does not provide accurate hotspot detection. For example, a hotspot typically causes an increase in packet loss but packet loss alone does not represent the existence of a hotspot because packet loss can result from wireless channel error, connectivity loss, etc. Therefore, accurate hotspot identification requires use of all the aforementioned indicators and only the combination of these indicators can correctly identify a hotspot. More importantly, the hotspot mitigation endeavor should not impose substantial signaling overhead nor compromise network connectivity.

1.3.3 Efficient Routing for Sensor Networks

Distributed wireless sensor networks are expected to have widespread applications within the coming decades, including tracking, monitoring, and emergency response systems for military and environmental purposes. These networks must be capable of adapting to changing environments and requirements. A sensor network application may need to alter its behavior to manage limited resources more efficiently, recover from broken network links, or change its functional behavior in response to commands issued by an operator. Since sensor devices are typically small in size and equipped with limited energy, the primary focus in the sensor networking community has been on energy conservation. Hence, various energy-conserving algorithms have been proposed for sensor networks that increase longevity of the network in exchange for reduced fidelity and increased latency.

The dynamic and lossy nature of wireless communication poses major challenges to reliable, self-organizing multi-hop networks. These non-ideal characteristics are more problematic with the primitive, low-power radio transceivers found in sensor networks, and raise fidelity issues that must be addressed. Therefore, the support for adequate fidelity is tightly coupled to link-quality and the routing decisions over these wireless links. As in a mobile ad hoc network, the fidelity of a sensor network depends on various network dynamics. Being equipped with unsophisticated transceivers, these network dynamics greatly impact the network performance and often deliver only a fraction of transported information.

In addition, sensor networks often exhibit non-isotropic radio ranges and possess asymmetric and unidirectional links. Therefore, the support for higher fidelity judiciously favors the bidirectional link over the unidirectional link in the path establishment phase. The link-layer reliability between two sensors can be perceived through mutual eavesdropping of transmissions or signaling exchanges.

However, efforts for enhanced fidelity should not impose substantial overhead. Since sensor networks may be idle for most of the time and active for a short period, route maintenance or link-state evaluations should not incur substantial overhead. While achieving a good routing path is very important, it is also crucial that a good path is attained at a reasonable cost.

1.4 Thesis Outline

In order to overcome the technical issues presented above, we propose the use of combination of analysis, simulations, and experimentation to best understand the problems and solution space. The outline of our study is as follows.

In Chapter 2, we present the design, implementation, and evaluation of the INSIGNIA QOS framework that is capable of supporting adaptive services in mobile ad hoc networks. The architecture includes a novel in-band signaling system that is lightweight in nature and highly responsive to network dynamics. The INSIGNIA signaling system is capable of quickly establishing, restoring, and adapting flows to meet time-varying resource availability. We introduce the notion of soft-state for the management of wireless resources. This approach is very effective in support of adaptive services and changing network topology. We also show that the adaptive mobile soft-state promotes better network utilization and resolves the problems of false restoration and resource lock-up found in soft-state driven mobile ad hoc networks.

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 mobile ad hoc networks. The INSIGNIA framework supports the following design features: (1) service differentiation and application adaptation, (2) fast and responsive in-band signaling in support of fast reservation and restoration, (3) distributed resource control using ‘soft-state’ resource management, (4) separation between routing, signaling and packet forwarding, and (5) 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 mobile ad hoc network.

In Chapter 3, we present extensive performance evaluation of the INSIGNIA system with AODV [30], DSR [31], and TORA [32]. We show that INSIGNIA improves performance for UDP and various TCP protocols (i.e., TCP-Reno [33], TCP-SACK [34], and TCP-Vegas [33]) under various node mobility and network loading conditions. We also evaluate the INSIGNIA system with Explicit Link Failure Notification (ELFN) [35], which is specifically designed to enhance TCP in mobile ad hoc networks.

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. We show that INSIGNIA is responsive to the mobility of nodes, load on the network and ability of applications to adapt. We believe that INSIGNIA is well suited to support adaptive real-time applications in mobile ad hoc networks.

In Chapter 4, we address the issues related to the congestion condition of wireless ad hoc networks called a hotspot. Hotspots represent transient but highly congested regions in wireless ad hoc networks that result in various networking problems. We demonstrate that hotspots exist even in lightly loaded mobile ad hoc networks and their existence can severely limit the performance. We present a simple protocol called Hotspot Mitigation Protocol (HMP) that works with existing best effort routing protocols to mitigate hotspots in wireless ad hoc networks. HMP tackles the problem right at the point of congestion, as opposed to traditional end-to-end approaches found in the literature. We show that traditional remedies such as end-to-end congestion control are often not effective in ad hoc networks and can limit the utilization and connectivity of the wireless network in the face of hotspots.

HMP is evaluated using both on-demand and proactive MANET routing protocols. HMP provides significant increases in network performance, improves network connectivity, and lowers routing overhead for on-demand routing protocols. In the case of proactive routing schemes, HMP provides some performance boosts for DSDV [36] but has limited success with the OLSR [37] protocol due to its design of routing packets through specially designated nodes. To get some hands-on experience with the protocol we also implemented HMP with AODV in a small-scale wireless testbed and confirmed the performance benefits observed under simulation. Based on our results, we recommend that future mobile ad hoc routing algorithms should incorporate the notion of hotspots into their protocols.

In Chapter 5, we investigate the classical issue of efficient/effective routing in the realm of the sensor network. It is observed that current routing algorithms for sensor networks provide poor information delivery (i.e., low fidelity), poor efficiency, and poor adaptability to changes in network condition. Since sensor devices are inherently limited in lifetime, expiring sensor devices together with other network dynamics continuously alter the network topology. We also anticipate that new sensors will be added to an existing network to extend the networking coverage, to improve sensing resolution, or to replace expiring sensor devices. Consequently, the topology of a sensor network changes continuously and the on-going information delivery is persistently afflicted by network dynamics. To tackle this important issue, we propose a new routing algorithm called solicitation-based forwarding (SOFA). The on-demand nature of SOFA makes it cost effective, its agile self-adapting capability makes it resilient to network dynamics, and its timely forwarding decision through solicitation-based handshakes makes it effective in packet delivery.

Our experimental testbed results confirm that SOFA provides excellent path convergence and supports responsive adaptations to network dynamics. In response to dynamic path conditions, SOFA quickly acquires alternative paths with minimal control overhead. Path maintenance in SOFA is minimal and when a new sensor is added to the network, it is integrated quickly and seamlessly. SOFA is also capable of integrating clusters of new sensors without incurring lengthy settling time. The combination of these attributes allows SOFA to provide improved fidelity over the baseline network. SOFA attempts to provide good paths between sources and the sink through series of hop-by-hop decisions. Each per-hop decision reflects the local condition of a node, so that the complete path represents a fusion of piecewise best forwarding nodes. Therefore, SOFA is most useful where network dynamics and topology changes are present. We claim that our proposal is appropriate for an event-driven sensor network because it is light-weight, maintenance-free when in idle state, provides fast convergence, responsively adapts to network topology changes, and integrates new sensors without delay.

1.5 Thesis Contribution

The contribution of this thesis can be summarized as follows:

1. In Chapter 2, we present the design, implementation, and evaluation of *INSIGNIA*, an IP-based quality of service framework that supports adaptive services in mobile ad hoc networks. The INSIGNIA framework is based on in-band signaling, soft-state management, and per-flow state management techniques. We show that in-band signaling is more suitable than explicit out-of-band approaches and soft-state management is well suited to support end-to-end quality of service in highly dynamic environments such as mobile ad hoc networks where network topology, node connectivity and end-to-end quality of service are strongly time-varying. INSIGNIA is designed to support fast reservation, fast restoration, and adaptation algorithms that help to counter time-varying network dynamics. To best of our knowledge, INSIGNIA is one of the first work that raised QOS issues in MANET and the first QOS signaling protocol designed specifically for MANETs.
2. In Chapter 3, we present extensive performance evaluation of INSIGNIA. We show that INSIGNIA is a general-purpose approach to deliver quality of service in mobile ad hoc network and provides operational transparency to a number of mobile ad hoc network routing protocols such as Ad hoc On-demand Distance Vector (AODV), Dynamic Source Routing (DSR) and the Temporally Ordered Routing Algorithm (TORA). We show that INSIGNIA provides substantial performance gains for UDP and TCP traffic under diverse networking conditions.

3. Hotspots represent transient but highly congested regions in mobile ad hoc networks that typically manifest increased packet loss, end-to-end delay, and out-of-order packets delivery. In Chapter 4, we present a simple, effective, and scalable *hotspot mitigation protocol* (HMP) where mobile nodes independently monitor local buffer occupancy, packet loss, and MAC contention and delay conditions, and take local actions in response to the emergence of hotspots, such as suppressing new route requests and rate controlling TCP flows. HMP is one of the first papers that identified the hotspots and it presents a simple but elegant solution to this vexing problem. We use analysis, simulations, and an experimental testbed to demonstrate that HMP effectively mitigates hotspots in mobile ad hoc networks. Moreover, HMP balances resource consumption among neighboring nodes, improves general network performance (i.e., end-to-end throughput, delay, packet loss, etc.), and improve the network connectivity by preventing premature network partitions.
4. In Chapter 5, we present the Solicitation-based Forwarding Algorithm (SOFA). Providing satisfactory fidelity for a sensor network is intrinsically challenging. Unpredictable network dynamics and the presence of transitional regions [85] in sensor networks significantly impact information delivery (i.e., fidelity). Link quality between multi-hop wireless sensor devices is highly unpredictable and often delivers only a fraction of the intended packets, even under the best conditions. Consequently, the amount of delivered information often fails to meet the fidelity requirement of an application. We argue that one of the major reasons for low-fidelity is due to non-responsive forwarding mechanisms commonly implemented in

existing sensor networks. Many existing routing protocols base their forwarding decisions on some form of statistical records of past communications or the quality of past beacon signals received. This approach fails to capture the link conditions at the exact time of forwarding packets across the link, limiting the effectiveness of forwarding decision. Therefore, we argue that forwarding decisions in sensor networks should be based on the actual network condition at the time of communication. We propose *solicitation-based forwarding* (SOFA), an agile, cost-effective, maintenance-free, and high-fidelity yielding hop-by-hop routing protocol that makes use of solicitation-based handshakes between a sender and multiple potential receivers at each hop to negotiate an appropriate forwarding path to a targeted destination (i.e., sink). We present the detailed design, implementation, and experimental evaluation of SOFA in a 36-node Mica-2 testbed using TinyOS.

Chapter 2

INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks

2.1. Introduction

Mobile ad hoc networks are autonomous distributed systems that comprise a number of mobile nodes connected by wireless links forming arbitrary time-varying wireless network topologies. Mobile nodes function as hosts and routers. As hosts, they represent source and destination nodes in the network while as routers, they represent intermediate nodes between a source and destination, providing store-and-forward services to neighboring nodes. Nodes that constitute the wireless network infrastructure are free to move randomly and organize themselves in arbitrary fashions. Therefore the wireless topology that interconnects mobile hosts/routers can change rapidly in unpredictable ways or remain relatively static over long periods of time. These bandwidth-constrained multi-hop networks typically support best effort voice and data communications where the achieved “goodput” is often lower than the maximum radio transmission rate after encountering the effects of multiple access, fading, noise, and interference, etc. In addition to being bandwidth constrained, mobile ad hoc networks are power constrained because network nodes rely on battery power for energy. Providing suitable quality of service (QOS) support for the delivery of real-time audio, video and data in mobile ad hoc networks presents a number of significant technical challenges.

Mobile ad hoc networks may be large, which makes the problem of network control difficult. The end-to-end communication abstraction between two communicating mobile hosts can be viewed as a complex “end-to-end channel” that may change route over time. There may be a number of possible routes between two communicating hosts over which data can flow, and each path may have different available capacity that may or may not meet the quality of service requirements of the desired service. Even if the selected path between a source-destination pair meets the user's needs at the session set-up time, the capacity and error characteristics observed along the path are likely to be time varying due to the multiple dynamics that operate in the network.

The fading effects resulting from host mobility cannot always be masked by the link layer and typically result in discernible effects on the application's perceptible quality (e.g., assured delivery of audio/video may degrade rapidly). This affects the capacity of a given path through the network, where links tend to degrade slowly at first and then rapidly drop out. This results in topological dynamics that operate on slower time scales than channel fades and other such discontinuities. Reacting to these network capacity dynamics over the appropriate time scale requires fast, lightweight, and responsive protocol operations. Flows must be established, maintained, and removed in mobile ad hoc networks over the course of a user-to-user session. Typically, “connections” (i.e., the establishment of “state” information at nodes along the path) need to be maintained and automatically renegotiated in response to the network topology dynamics and link quality changes. Since resources are scarce in these networks, any protocol signaling overhead needed to maintain connections limits the utilization of the network. Therefore, bandwidth required to support signaling systems must be kept to a minimum. This places emphasis on minimizing the signaling required to establish, maintain, restore, and tear down network states associated with user sessions. In addition, due to the disconnected nature of maintaining state in mobile ad hoc networks, explicit tear-down mechanisms (e.g., disconnect signaling) are impractical. This is due to the fact that it is infeasible to explicitly remove network state (established during session setup) in portions of the network that are out of radio contact of a signaling controller due to topology changes.

There is a need for new mobile ad hoc architectures, services, and protocols to be developed in response to these challenges. New control systems need to be highly adaptive and responsive to changes in the available resources along the path between two communicating mobile hosts. Future protocols need to be capable of differentiating between the different service requirements of user sessions (e.g., continuous media flows, microflows, RPC, etc.). Packets associated with a flow traversing intermediate nodes (as illustrated in Figure 2-1) between a source and destination may, for example, require special processing to meet end-to-end bandwidth and delay constraints. When building quality of service support into mobile ad hoc networks the design of fast routing algorithms that can efficiently track network topology-changes is important. Mobile ad hoc network routing protocols need to work in unison with efficient signaling, control, and management mechanisms to achieve end-to-end service quality. These mechanisms should consume minimal bandwidth in operation and react promptly to changes in the network state (viewed in terms of changes in the network topology) and flow state (viewed in terms of changes in the observed end-to-end quality of service).

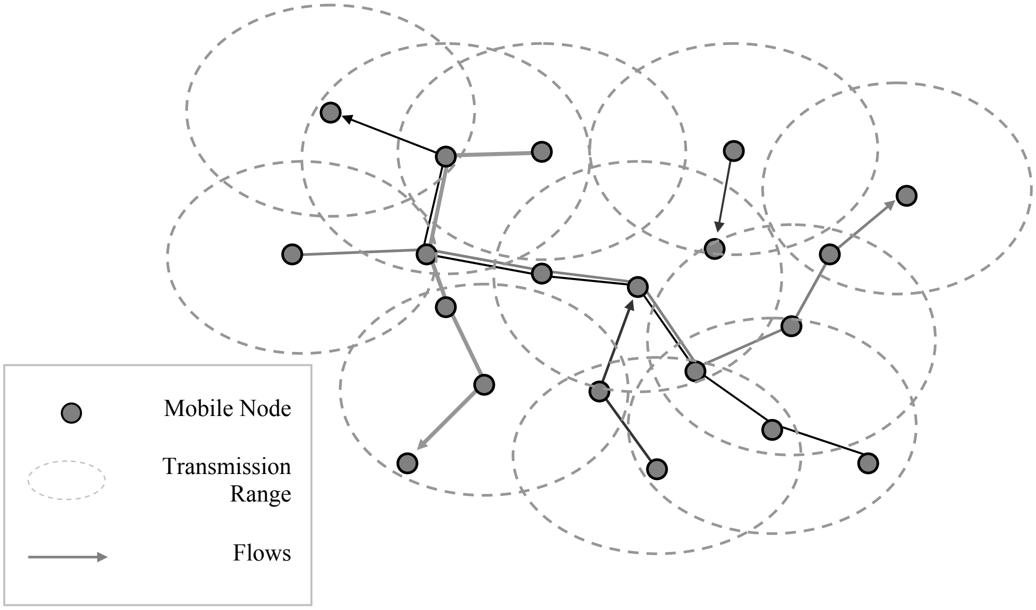

Figure 2-1: Mobile Ad Hoc Networking

In this chapter, we present the design, implementation, and evaluation of the INSIGNIA QOS Framework that supports the delivery of adaptive services in mobile ad hoc networks. A key component of our QOS framework is the INSIGNIA signaling system, an in-band signaling system that supports fast reservation, restoration, and adaptation algorithms that are specifically designed to deliver adaptive service. The signaling system is designed to be lightweight and highly responsive to changes in network topology, node connectivity, and end-to-end quality of service conditions.

The structure of this chapter is as follows. We discuss our framework in the context of the related work and present the main design considerations that have influenced our thinking in Sections 2.2 and 2.3, respectively. Section 2.4 presents an overview of the INSIGNIA QOS framework. The detailed design of the INSIGNIA signaling system is given in Section 2.5. We evaluate our QOS framework in Section 2.6, paying particular attention to the performance of the signaling system under a variety of network conditions. Our simulation results show the benefit of the INSIGNIA QOS framework under diverse mobility, traffic, and channel conditions in support of fast reservation, restoration, and adaptation. Finally, we present our conclusion in Section 2.7.

2.2 Related Work

Research and development of mobile ad hoc networking technology is proceeding in both academia and industry under military and commercial sponsorship. Current military research projects such as the Army Research Office Focused Research

Initiatives, the Army Research Laboratory Federated Laboratory, and the DARPA Global Mobile Information Systems (GloMo) program [59] are producing new technologies.

There has been little research in the area of supporting quality of service in mobile ad hoc networks, however. What work exists tends to be based on distributed scheduling algorithms [72] that address rescheduling when the network topology changes and QOS-based medium access controllers [70]. Typically, these schemes are based on a single link layer network technology and not on an interconnection of different wireless technologies at the IP layer. In addition, the work does not address suitable support for adaptive QOS paradigms that are required to deliver adaptive services in mobile ad hoc networks. In this chapter, we propose an IP-based QOS framework, adaptive services, and support protocols incorporating a soft-state [43] resource management system. This system is based on in-band signaling techniques supporting reservation across multiple link layer radio technologies that map to specific link layer access technologies for distributed packet scheduling. Our contribution addresses a suitable IP level control architecture for delivering adaptive services in mobile ad hoc networks. We do not, however, propose any new distributed scheduling techniques. Rather, we leverage the existing body of work found in the literature as a basis for the provision of QOS support over radios.

In [48] [49], multi-hop, multi-cluster packet radio network architectures are proposed. The provisioning of quality of service is discussed based on "dynamic virtual circuit" communications derived from wireline network control and signaling found in

ATM networks. This approach relies on a “circuit” model that requires explicit connection management and the establishment of hard state in the network prior to communication. We believe there is a need to investigate alternative network models that are more responsive to the dynamics found in ad hoc networks other than the hard-state virtual circuits. Typically, virtual circuits are established across mobile ad hoc networks using explicit “out-of-band” signaling to set up reservations for the duration of the call/session holding time. We believe that flows/sessions should be established and maintained using a faster, more responsive system based on soft-state and in-band signaling paradigms. We believe that virtual circuits lack the intrinsic flexibility needed to adapt to the dynamics found in mobile ad hoc networks and that the notion of “soft-state connections” driven by in-band techniques is more suitable. There is a need to develop new QOS architectures that can provide fast reservation, responsive restoration, and seamless adaptation to mobile ad hoc network dynamics based on the inherent flexibility, robustness, and scalability found in IP networks.

Delivering end-to-end service quality in mobile ad hoc networks is intrinsically linked to the performance of the routing protocol because new routes or alternative routes between source-destination pairs need to be periodically computed during ongoing sessions. The IETF Mobile Ad hoc Networks (MANET) Working Group [56] recently began to standardize inter-network layer technologies (i.e., routing and support protocols). As such, it is presently focused on standardizing network-layer routing protocols suitable for supporting best effort packet delivery in IP-based networks. Within this context there have been a number of proposals for efficient routing that dynamically track changes in mobile ad hoc network topology including the Temporally Ordered Routing Algorithm (TORA) [32], Dynamic Source Routing [31], Zone Routing Protocol [41] and Ad Hoc On demand Distance Vector Routing Protocol [30]. The performance of a QOS framework will rely on the speed at which routing protocols can compute new routes (if no alternative route is currently cached) after topology changes have occurred. The delay in computing new routes will have an impact on the QOS delivered to on-going sessions. For a comparison of mobile ad hoc routing protocols see [45].

2.3. Design Considerations

2.3.1 Adaptive Services

The most suitable service paradigm for mobile ad hoc networks is adaptive in nature. We observe that adaptive voice and video applications operating in mobile cellular networks are capable of responding to packet loss, delay jitter, changes in available bandwidth, and handoff while maintaining some level of service quality [46]. While adaptive multimedia applications can respond to network dynamics they typically require some minimum bandwidth assurance below which they are rendered useless.

The INSIGNIA QOS framework is designed to support adaptive services as a primary goal. In this context, adaptive services provide minimum bandwidth assurances to real-time voice and video flows and data allowing for enhanced levels (i.e., maximum bandwidth) of service to be delivered when resources become available.

A flow represents a sequence of packets sent from a single source to one or more destinations representing a single media type (e.g., voice, video, etc.). Flows require admission control, resource reservation, and maintenance at all intermediate routers between a source and destination to provide end-to-end quality of service support. Typically, continuous media flows are long lived in comparison to microflows, which represent short-lived flows (e.g., web style client/server interactions) that comprise a limited train of data packets. We use the terms “session,” “flow,” “continuous media flow,” and “microflow” interchangeably in this chapter. The INSIGNIA QOS framework is designed to transparently support the requirements of continuous media flows and microflows. Adaptive services support applications that require *base QOS* (i.e., minimum bandwidth) and *enhanced QOS* (i.e., maximum bandwidth) assurances, respectively. The semantics of the adaptive service provides preference to packets associated with the base QOS over enhanced QOS. Adaptation is an application-specific process. Some applications may be incapable of adapting while others may adapt discretely (e.g., scalable profiles of MPEG2) or continuously (e.g., dynamic rate-shaped applications [46]). The time scale over which applications can adapt is also application specific. For example, greedy data applications (e.g., image downloads) may want to take advantage of any change in available bandwidth at any time. In contrast, adaptive continuous media applications (e.g., audio and video) may prefer to follow trends (via some low pass filtering scheme) in available bandwidth based on slower adaptation time scales, preferring some level of “stable” service delivery rather than responding to every instantaneous change in bandwidth availability. Adaptive applications therefore should manage the adaptation process and dictate the time scales and semantics of their adaptation process. Given this observation, our QOS framework is designed to adapt user sessions to the available level of service without explicit signaling between source-destination pairs. In this case the network and application adapt to different dynamics. The network adapts (via restoration algorithms) to changes in topology and measured channel conditions while trying to deliver base and enhanced QOS. Applications adapt to the observed end-to-end QOS fluctuations within the prescribed max-min limits based on application specific adaptation time scales. This observation drives a number of architectural design decisions.

2.3.2 Separation of Routing, Signaling and Forwarding

There has been a growing amount of work in the area of QOS routing for fixed networks. Here the routing protocols interact with resource management to establish paths through the network that meet end-to-end QOS requirements (i.e., delay, bandwidth, possibly multi-metrics demands). In this case there is a certain level of integration of resource management and routing. One could apply such an approach to MANET routing protocols given that the time scales over which new routes are computed are much faster than traditionally found in the case of routing in fixed infrastructures. While we believe this a promising approach (see the CEDAR [47] proposal) we note that the time scales over which session setup and routing (i.e., computing new routes) operate are distinct and functionally independent tasks.

Therefore, we believe that signaling, resource management, and routing should be modeled independently in the network architecture.

We consider that MANET routing protocols should not be burdened with the integration of QOS functionality that may be tailored toward specific QOS models. Rather, we argue that it is better to maintain a clean separation between routing, signaling, and forwarding. These architectural components are rather different from one another in the algorithms they implement and in the time scales over which they operate. Our approach is to develop a QOS framework that can ``pluggin" a wide variety of routing protocols. In this case, resource reservation and signaling will be capable of interacting with any number of routing protocols to provide end-to-end QOS support. Different MANET routing protocols clearly perform differently [45] in response to topology changes while the QOS framework attempts to maintain end-to-end service quality.

2.3.3 In-Band Signaling

In-band signaling systems are capable of operating close to packet transmission speeds and are therefore well suited toward responding to fast time scale dynamics found in mobile ad hoc environments, as illustrated in Figure 2-1. The term “in-band signaling” refers to the fact that the control information is carried along with data. In contrast, out-of-band signaling systems (e.g. Internet’s RSVP, ATM’s UNI, etc.) are incapable of responding to such fast time-scale dynamics because out-of-band signaling systems require maintenance of source route information and respond to topology changes by directly signaling “affected mobiles” to allocate/free resources. In some cases, this is impossible to do due to lack of connectivity between the “affected router” and the signaling entity that attempts to deallocate resources over the old path.

The term “out-of-band signaling” refers to fact that the control information is typically carried in separate control packets and on channels that may be distinct from the data path. Based on an in-band approach, the INSIGNIA signaling system can restore flowstate (i.e., a reservation) in response to topology changes within the interval of two consecutive IP packets under ideal conditions. INSIGNIA performance relies on the speed at which the routing protocol can re-compute new routes if no alternative route is cached after topology changes. Out-of-band signaling systems, for example, would need to maintain source route information and respond to topology changes by directly signaling intermediate routers on an old path to allocate/free radio resources. In many case, this is impossible to do if the affected router is out of radio contact from the signaling entity that attempts to de-allocate resources over the old path.

2.3.4 Soft-State Management

Maintaining the QOS of adaptive flows in mobile ad hoc networks is one of the most challenging aspects of the INSIGNIA QOS framework. In wireline networks that support quality of service and state management, the route and the reservation between source-destination pairs remain fixed for the duration of a session. This style of hard-state connection oriented communications (e.g., virtual circuit) guarantees quality of service for the duration of the session holding time. However, these techniques are not flexible enough in mobile ad hoc networks, where the path and reservation need to dynamically respond to topology changes in a timely manner.

We believe that a soft-state approach to state management at intermediate routing nodes is suitable for the management of reservations in mobile ad hoc networks. Such an approach models the transient nature of network reservations, which have to be responsive to fast time-scale wireless dynamics, moderate time-scale mobility changes and longer time scale session "holding times." Based on the work by Clark [43], soft-state relies on the fact that a source sends data packets along an existing path. If a data packet arrives at a mobile router and no reservation exists then admission control and resource reservations attempt to establish soft-state. Subsequent reception of data packets (associated with a reservation) at that router are used to refresh the existing soft-state reservation. This is called a "soft-connection" when considered on an end-to-end basis and in relation to the virtual circuit hard-state model. When an intermediate node receives a data packet that has an existing reservation it reconfirms the reservation over the next interval. Therefore the holding time for a soft connection is based on the soft-state timer interval and not based on session duration holding time. If a new packet is not received within the soft-state timer interval then resources are released and flow states removed in a fully decentralized manner.

We believe that the development of new QOS frameworks based on the notion of in-band signaling and soft-state management and constructed with separation of routing, QOS, signaling and forwarding functions will provide a responsive, scalable and flexible solution for delivering adaptive services in mobile ad hoc networks.

2.4. The INSIGNIA QOS Framework

The INSIGNIA QOS framework allows packet audio, video and real-time data applications to specify their maximum and minimum bandwidth needs and plays a central role in resource allocation, restoration control and session adaptation between communicating mobile hosts. Based on availability of end-to-end bandwidth, QOS mechanisms attempt to provide assurances in support of adaptive services. To support adaptive service, the INSIGNIA QOS framework establishes and maintains reservations for continuous media flows and micro-flows. To support these communication services the INSIGNIA QOS framework comprises the following architectural components as illustrated in Figure 2-2:

- *In-band signaling* establishes, restores, adapts and tears down adaptive services between source-destination pairs. Flow restoration algorithms respond to dynamic route changes and adaptation algorithms respond to changes in available bandwidth. Based on an in-band signaling approach that explicitly carries control information in the IP packet header, flows/sessions can be rapidly established, restored, adapted and released in response to wireless impairments and topology changes.

- *Admission control* is responsible for allocating bandwidth to flows based on the maximum/minimum bandwidth (i.e., base and enhanced QOS) requested. Once resources have been allocated they are periodically refreshed by a soft-state mechanism through the reception of data packets. Admission control testing is based on the measured channel capacity/utilization and requested bandwidth. To keep the signaling protocol simple and lightweight, new reservation requests do not impact existing reservations.

- *Packet forwarding* classifies incoming packets and forwards them to the appropriate module (viz. routing, signaling, local applications, packet scheduling modules). Signaling messages are processed by INSIGNIA signaling, and data packets are delivered locally (as illustrated by the dashed line in Figure 2-2) or forwarded to the packet scheduling module (as illustrated by the bold line in Figure 2-2) for transmission on to the next hop.

- *Routing protocol* dynamically tracks changes in ad hoc network topology making the routing table visible to the node's packet forwarding engine. The QOS framework assumes the availability of a generic set of MANET routing protocols [42] that can be plugged into the architecture. The QOS framework assumes that the routing protocol provides new routes, either proactively or on-demand, in the case of topology changes.

- *Packet scheduling* responds to location-dependent channel conditions when scheduling packets in wireless networks [62]. A wide variety of scheduling disciplines can be used to realize the packet scheduling module and the service model. Currently, we have implemented a weighted round robin [62] [64] service discipline based on an implementation [69] of deficit round robin that has been extended to provide compensation in the case of location dependent channel conditions between mobile nodes.

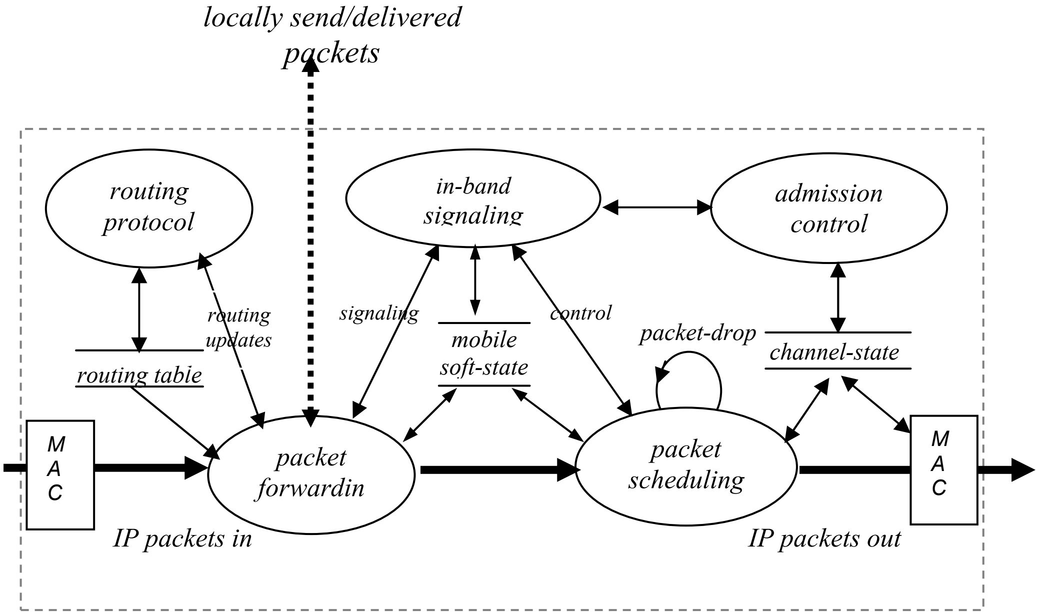

Figure 2-2: INSIGNIA QOS Framework

- *Medium access control (MAC)* provides quality of service driven access to the shared wireless media for adaptive wireless and best effort services. The INSIGNIA QOS framework is designed to be transparent to any underlying media access control protocols and is positioned to operate over multiple link layer technologies at the IP layer. However, the performance of the framework is strongly coupled to the provisioning of QOS support provided by specific medium access controllers.

2.5. The INSIGNIA Signaling System

The INSIGNIA signaling system plays an important role in establishing, adapting, restoring, and terminating end-to-end reservations. In what follows, we describe the INSIGNIA in-band signaling approach. The signaling system is designed to be lightweight in terms of the amount of bandwidth consumed for network control and to be capable of reacting to fast network dynamics such as rapid host mobility, wireless link degradation, and intermittent session connectivity. We discuss the protocol command and then the protocol mechanisms.

2.5.1 Protocol Commands

Protocol commands are encoded using the IP option field and include service mode, payload type, bandwidth indicator and bandwidth request field as illustrated in Figure 2-3. By adopting an INSIGNIA IP option in each IP packet header the complexity of supporting packet encapsulation inside the network is avoided. These protocol commands supports the signaling algorithms discussed in Section 2.5.2 including flow reservation, restoration, and adaptation mechanisms. The protocol commands drive the state operations of the protocol. Figure 2-4 presents a simplified view of the finite state machines for a source host, intermediate router, and destination host. These three state machines capture the major event/actions and resulting state transitions. We use these state machines to illustrate the dynamics of the INSIGNIA signaling system.

2.5.1.1 Service Mode

When a source node wants to establish a fast reservation to a destination node it sets the *reservation (RES) mode* bit in the INSIGNIA IP option service mode of a data packet and send the packet toward the destination. On reception of a RES packet intermediate routing nodes execute admission control to accept or deny the request. When a node accepts a request, resources are committed and subsequent packets are scheduled accordingly. In contrast, if the reservation is denied, packets are treated as *best effort (BE) mode* packets.

In the case where a RES packet is received and no resources have been allocated, the admission controller attempts to make a new reservation. This condition commonly occurs when flows are rerouted during the lifetime of an ongoing session due to host mobility. When the destination receives a RES packet it sends a QOS report to the source node indicating that an end-to-end reservation has been established and transitions its internal state from best effort to reservation state as illustrated in Figure 2-4(c).

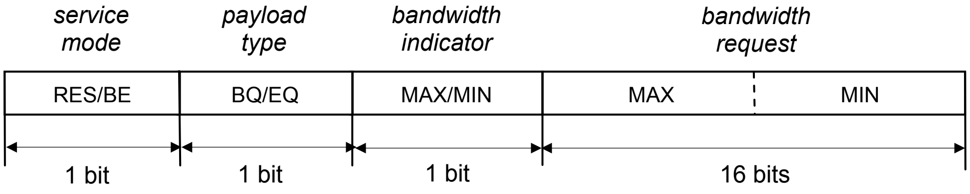

Figure 2-3: INSIGNIA IP Option

The service mode indicates the level of service assurance requested in support of the adaptive services. The interpretation of the service mode, which indicates a RES or BE packet, is dependent on the payload type and bandwidth indicator discussed in Section 2.5.1.3 and Section 2.5.1.4, respectively. A packet with the service mode set to RES and bandwidth indictor set to *MAX* or *MIN* is attempting to set-up a max-reserved or min-reserved service, respectively. The bandwidth requirements of the flow are carried in the bandwidth request field, as illustrated in Figure 2-3. A RES packet may be degraded to BE service in the case of rerouting or insufficient resources availability along the new/existing route. Note that a BE packet requires no resource reservation to be made.

The IP option also carries an indication of the payload type, which identifies whether the packet is a base QOS *(BQ)* or enhanced QOS *(EQ)* packet as discussed in Section 2.5.1.3. Using the “packet state” (service mode/payload type/bandwidth indicator) one can determine what component of the flow is degraded. Reception of a BE/EQ/MIN packet or RES/BQ/MIN indicates that the enhanced QOS packets have been degraded to best effort service. By monitoring the packet state the destination node can issue scaling/drop commands to the source based on the destination state machine illustrated in Figure 2-4(c).

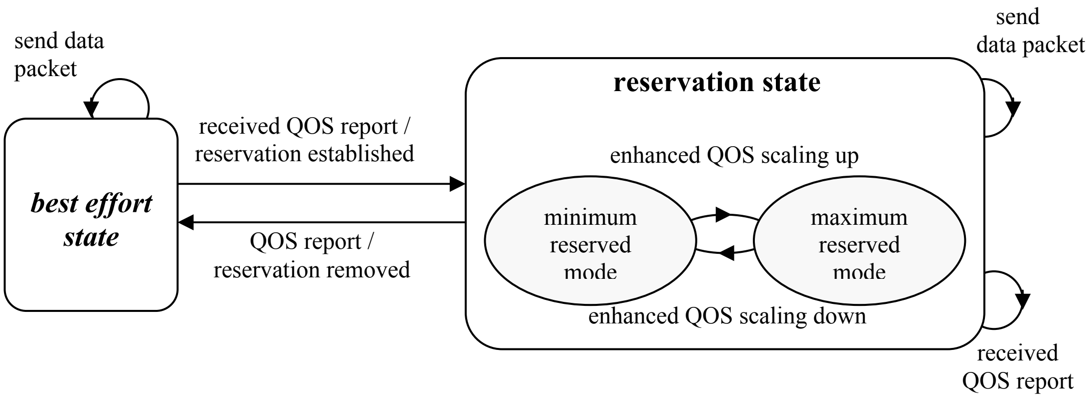

Figure 2-4(a). State Machine at a Source Mobile Host

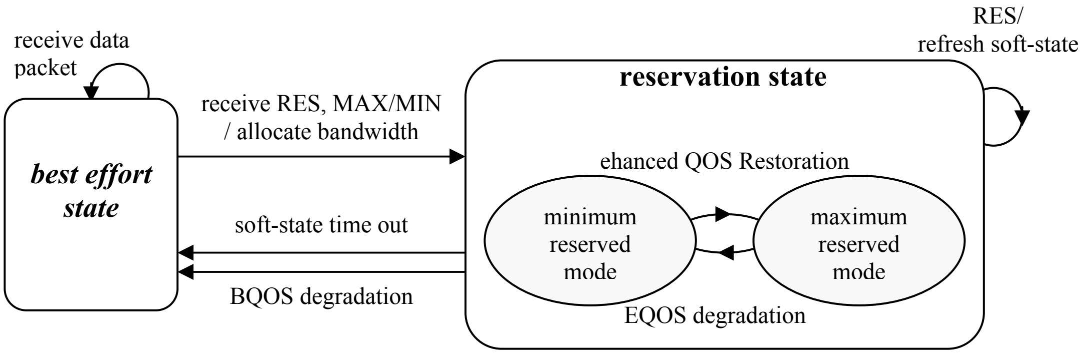

Figure 2-4(b). State Machine at an Intermediate Mobile Node

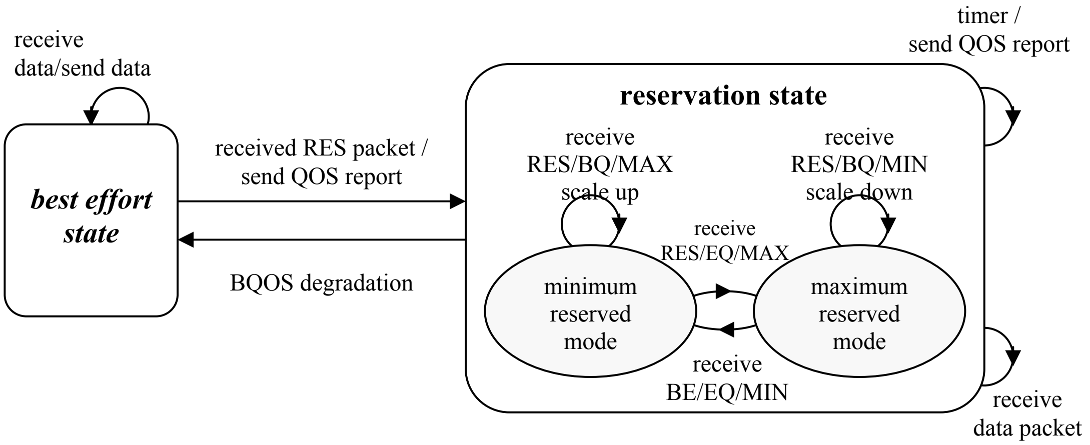

Figure 2-4(c). State Machine at Destination Mobile Host

As shown in Figure 2-4 the source, intermediate and destination state machines support two reservation sub-states:

- *max-reserved mode* provides reservation for a flow's base QOS and enhanced QOS packets. This type of service requires successful end-to-end reservation to meet a flow's maximum bandwidth needs (e.g., RES/EQ/MAX).
- *min-reserved mode* provides reservation for the base QOS and best effort delivery for the enhanced QOS components (if it exists). This service mode typically occurs when max-reserved flows experience degradation in the network. For example, max-reserved flows may encounter mobile nodes that lack resources to support both the base and enhanced QOS, resulting in the degradation of enhanced QOS packets to best effort delivery (e.g., BE/EQ/MIN).

2.5.1.2 Bandwidth Request

The bandwidth request allows a source to specify its maximum *(MAX*) and minimum *(MIN)* bandwidth requirements for adaptive services. This assumes that the source has selected the RES service mode. A source may also simply specify a minimum or a maximum bandwidth requirement. For adaptive services the base QOS (min-reserved service) is supported by the minimum bandwidth, whereas the maximum bandwidth supports the delivery of the base and enhanced QOS (max-reserved service) between source-destination pairs. Flows are represented as having minimum and maximum bandwidth requirements. This characterization is commonly used for multi-resolution traffic (e.g., MPEG audio and video), adaptive real-time data that has discrete max-min requirements, and differential services that support prioritization of aggregated data in the Internet.

2.5.1.3 Payload Type

The payload field indicates the type of packet being transported. INSIGNIA supports two types of payload called base QOS (BQ) and enhanced QOS (EQ), which are reserved via distributed end-to-end admission control and resource reservation. The semantics of the adaptive services are related to the payload type and available resources (e.g., enhanced QOS requires that maximum bandwidth requirements can be met along the path between a source-destination pair). The semantics of the base and enhanced QOS are applications specific. They can represent a simple prioritization scheme between packets, differential services, or self-contained packet streams associated with multi-resolution flows. The adaptation process may force adaptive flows to degrade when insufficient resources are available to support the maximum bandwidth along the existing path or during restoration when the new path has insufficient resources. For example, if there is only sufficient bandwidth to meet the minimum bandwidth requirement needs of the base QOS, enhanced QOS packets are degraded to best-effort packets at bottleneck nodes by simply flipping the service mode of EQ packets from RES to BE. When a down stream node detects degraded packets, they release any resources that may have previously allocated to support the transport of enhanced QOS packets. The adaptation process (discussed in Section 2.5.2.5) is also capable of scaling flows up by taking advantage of any of additional bandwidth availability that may be encountered along a new/existing path. In this case, a flow could be “scaled-up” from min-reserved to max-reserved mode delivery, as indicated in Figure 2-4(a) and 2-4(c).

2.5.1.4 Bandwidth Indicator

A bandwidth indicator plays an important role during reservation setup and adaptation. During reservation establishment the bandwidth indicator reflects the resource availability at intermediate nodes along the path between a source-destination pairs. Reception of a setup request packet with the bandwidth indicator bit set to MAX indicates that all nodes enroute have sufficient resources to support the maximum bandwidth requested (i.e., max-reserved mode). In contrast, a bandwidth indicator set to MIN implies that at least one of the intermediate nodes between the source and destination is a bottleneck node and insufficient bandwidth is available to meet the maximum bandwidth requirement; that is, only min-reserved mode delivery can be supported. In this case, adaptation algorithms at the destination can trigger the signaling protocol to release any over-allocated resources between the source and bottleneck node by issuing a “drop” command to the source node (see Section 2.5.2.5 on adaptation). A bandwidth indicator set to MIN does, however, indicate that the mobile ad hoc network can support the minimum requested bandwidth (i.e., min-reserved mode). The bandwidth indicator is also utilized during the adaptation of ongoing sessions in this manner. The adaptation mechanism resident at the destination host continuously monitors the bandwidth indicator to determine if the additional bandwidth is available to support better service quality.

2.5.2 Protocol Operations

In what follows, we provide an overview of the main protocol mechanisms and state machines for the source, intermediate router and destination nodes as illustrated in Figure 2-4. The key signaling components include reservation establishment, QOS reporting, soft-state management, flow restoration, and flow adaptation.

2.5.2.1 Fast Reservation

To establish adaptive flows, source nodes initiate reservations by setting the appropriate field in the IP option in data messages before forwarding “reservation request” packets on toward destination nodes. A reservation request packet is characterized as having the service mode set to RES, payload set to BQ/EQ and bandwidth indictor to MAX/MIN and valid bandwidth requirements. Reservation packets traverse intermediate nodes executing admission control modules, allocating resources, and establishing flow-state at all intermediate nodes between source-destination pairs, as illustrated in Figure 2-5. A source node continues to send reservation packets until the destination node completes the reservation setup phase by informing the source node of the status of the flow establishment phase using QOS reporting, as shown in Figure 2-6.

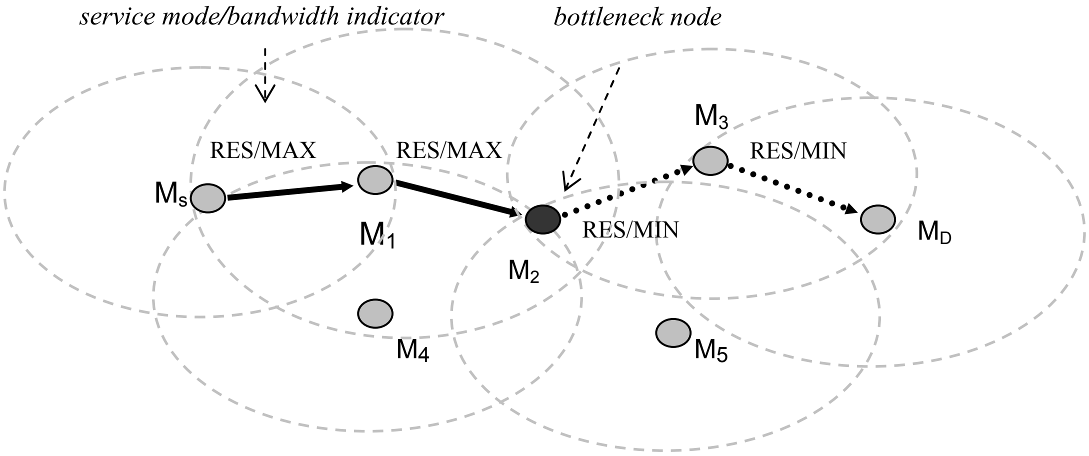

Figure 2-5: Adaptive Service Flow Reservation

The establishment of an adaptive flow is illustrated in Figure 2-5. A source node (M_S) requests maximum resource allocation and node M_1 performs admission control upon reception of the reservation packet. Resources are allocated if available, and the reservation packet is forwarded to the next node M_2. This process is repeated on a hop-by-hop basis until the reservation packet reaches the destination mobile M_D. The destination node determines the resource allocation status by checking the packet state (i.e., service mode, payload type, and bandwidth indicator). The QOS reporting mechanism is used to inform the source node of the reservation status enroute. As far as the destination node is concerned the reservation phase is complete on reception of the first RES packet. From the example shown in Figure 2-5, we see that only the minimum bandwidth is supported between M_2 and M_3 and subsequent nodes receiving the request packet avoid allocating resources for the maximum.

When a reservation is received at the destination node, the signaling module checks the flow establishment status. The status is determined by inspecting the IP option field service mode, which should be set to RES. If the bandwidth indication is set to MAX, this implies that all nodes between a source-destination pair have successfully allocated resources to meet the base and enhanced bandwidth needs in support off the max-reserved mode. On the other hand, if the bandwidth indication is set to MIN this indicates that only the base QOS can be currently supported (i.e., min-reserved mode). In this case, all reservation packets with a payload of EQ received at a destination will have their service level flipped from RES to BE by the bottleneck node. As a result "partial reservations" will exist between the source and bottleneck node (e.g., between M_S and M_2 in Figure 2-5). In the case of partial reservations, resources remain reserved between the source and the bottleneck node until explicitly released. Release of partial reserved resources can be initiated by the source based on feed back during the reservation phase or as part of the adaptation process where the destination can issue "scale-down/drop" commands to a source node. This will have the effect of clearing any partial reservation (e.g., between M_S and M_2 in Figure 2-5). An application may choose not to deallocate a partial reservation, hedging that bandwidth will become available at the bottleneck node allowing for a full end-to-end reservation to be made in due course.

Note that if a reservation has been established for the maximum reserved state and a RES/BQ/MIN packet is persistently received in this substate then the state machine determines that the enhanced QOS packets have been degraded and transitions to minimum reserved state in anticipation of scaling back up. This behavior is illustrated in Figure 2-4(c). Degradation of this sort can occur at intermediate node due to insufficient resources to support a new reservation, or an ongoing flow is degraded due to rerouting or insufficient resource availability on the new/existing path. The state information maintained at the destination can decode which of these conditions occurred.

2.5.2.2 QOS Reporting

QOS reporting is used to inform source nodes of the ongoing status of flows. Destination nodes actively monitor ongoing flows inspecting status information (e.g., bandwidth indication) and measuring the delivered QOS (e.g., packet loss, delay, throughput, etc.). QOS reports are also sent to source nodes for completing reservation phase and on a periodic basis for managing end-to-end adaptations. QOS reports do not have to travel on the reverse path toward the source. Typically they will take an alternate route through the ad hoc network as illustrated in Figure 2-6. Although the QOS reports are basically generated periodically according to the applications' sensitivity to the service quality, QOS reports are sent immediately when required (i.e., typically actions related to adaptation).

In the case where only the BQ packets can be supported, as is the case with the min-reserved mode, the signaling systems at the source "flips" the service mode of the BQ packets from RES to BE with all "degraded" packets sent as best effort. Any partial reservations that may exist between a source and destination nodes are automatically timed out after “flipping” the state variable in the EQ packets. Since there is a lack of EQ packets with the RES bit set at intermediate routers any associated resources are released (e.g., between M_S and M_2 in Figure 2-5) allowing other competing flows to contend for these resources. In a similar fashion QOS reports are also used as part of the ongoing adaptation process that responds to mobility and resources change in the mobile ad hoc network. The adaptation process is discussed in Section 2.5.2.5.

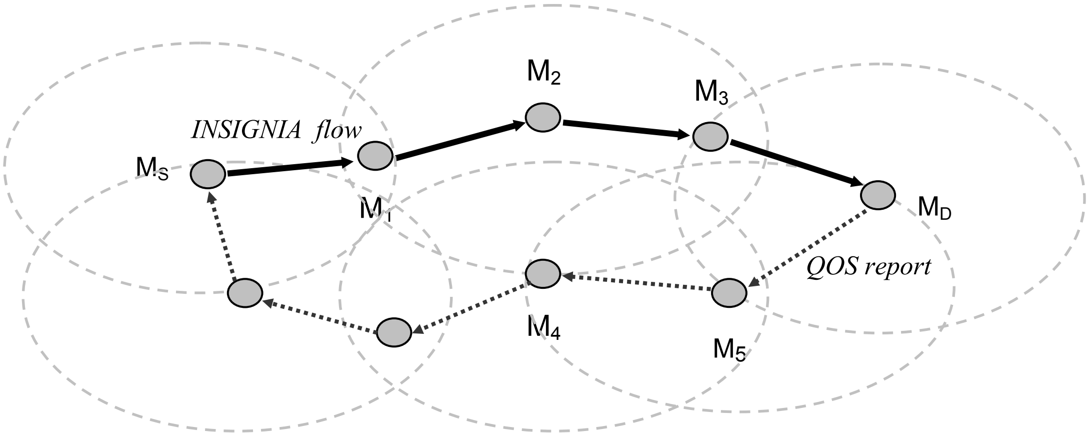

Figure 2-6: QOS Reporting

2.5.2.3 Soft-State Management

Reservations made at intermediate routing nodes between source and destination pairs are driven by soft-state management, as indicated by Figure 2-4(b). A soft-state approach is well suited for management of resources in dynamic environment, where the path and reservation associated with a flow may change rapidly. The transmission of data packets is strongly coupled to maintenance of flow states (i.e., reservations). In other words, as the route changes in the network, new reservations will be automatically restored by the restoration mechanism. A major benefit of soft-state is that resources allocated during flow establishment are automatically removed when the path changes. For example, the mobility of node M_2 in Figure 2-7 will cause flows to be rerouted to via intermediate routers M_1-M_4-M_3. Due to the absence of reserved mode data packets at node M_2 the node will automatically release resources associated with the flow without any interaction from any explicit controller.

Once admission control has accepted a request for a new flow soft-state management starts the soft-state timer associated with the new or rerouted flow. The soft-state timer is continually refreshed as long as packets associated with a flow are periodically received at intermediate routers. In contrast, if packets are not received (e.g., due to rerouting) then the soft state is not refreshed but times out with the result of deallocating any resources. Since data packets are used to maintain the state at intermediate nodes we couple the data rate of flows to the soft-state timer value. In Section 2.6.4, we evaluate the performance of a fixed and dynamic scheme for determining the soft-state timer value. The fixed scheme simply sets a value for all flows regardless of the data rate of individual flows (e.g., RSVP recommends 30 sec), and the dynamic scheme tracks the changing data rate of individual flows and sets the soft-state timer accordingly.

2.5.2.4 Restoration

Flows are often rerouted within the lifetime of ongoing sessions due to host mobility. The goal of flow restoration is to reestablish reservation as quickly and efficiently as possible. Rerouting active flows involves the routing protocol (to determine a new route), admission control, and resources reservation for nodes that belong to a new path. Restoration procedures also call for the removal of old flow state at nodes along the old path. In an ideal case, the restoration of flows 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." If no alternative route is cached, the performance of the restoration algorithm is coupled to the speed at which the routing protocols can discover a new path.

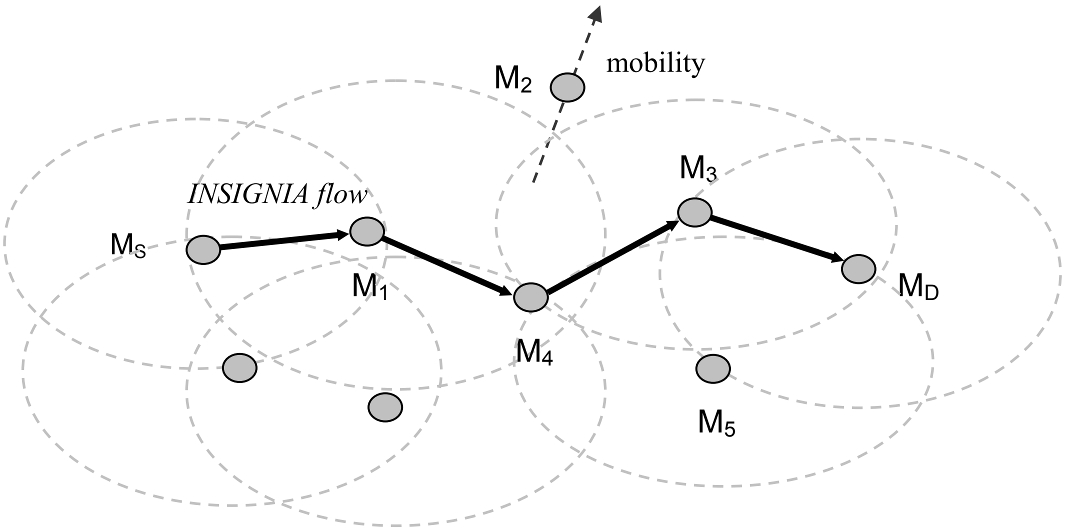

Figure 2-7. Rerouting and Restoration

As illustrated in Figure 2-7, network dynamics trigger rerouting and service degradation. In this example, mobile host M_2 moves out of radio contact and connectivity is lost in Figure 2-7. The forwarding router node, M_1 , interacts with the routing protocol and forwards packets along a new route. The signaling system at intermediate router M_4 receives packets and inspects its flow soft-state table. If a reservation does not exist for newly arriving packets then the signaling module invokes admission control and attempts to allocate resources for the flow. Note that when a rerouted packet arrives at node M_3 the forwarding engine detects that a reservation exists and treats the packet as any other packet with a reservation. In other words, the packets are routed back to the existing path, where a reservation is still present. Such scenarios are frequently observed in our experimental systems, discussed in Section 2.6, with the result of minimizing any service disruption due to rerouting. Soft-state timers ensure that the flow state is still intact at M_3 and that state along the old path (i.e., mobile host M_2) is removed in an efficient manner.

When an adaptive flow is rerouted to a 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 the flow. In this instance the state associated with a flow is timed out and resources are deallocated. A reservation may be restored if the resources free up at a bottleneck node (e.g., mobile node M_4 in Figure 2-8) or further rerouting may allow the restoration to complete. We call this type of restoration “degraded restoration.” A flow may remain degraded for the duration of the session and never be restored; this is described as "permanent degradation." The enhanced QOS component of an adaptive flow may be degraded to best effort service (i.e., min-reserved mode) during the flow restoration process if the nodes along the new path can only support the minimum bandwidth requirement. If the degradation of enhanced QOS packets persist, it may cause service disruption and trigger the destination mobile node to invoke its adaptation procedure to "scale down" or "drop" packets rather than live with degraded quality. Adaptation mechanisms located at destination nodes are capable of responding to changes in network resource availability through scale down, scale up, and drop actions in response to network conditions.

During the restoration process, the INSIGNIA framework does not favor rerouted flows over existing flows (e.g., by forcing existing flows to scale down to their minimum requirements to allow rerouted or new flows to be admitted). In this sense, INSIGNIA avoids the introduction of additional service fluctuations to existing flows in support of the restoration of rerouted flows. As a result of this policy, admission control simply rejects/scales down any rerouted flows when insufficient resources are available along a new path.

Three types of restoration are supported by the INSIGNIA QOS framework:

- An *immediate restoration* occurs when a rerouted flow immediately recovers its original reservation; that is, a max-reserved mode flow is immediately restored as a max-reserved mode flow and a min-reserved mode flow as a min-reserved mode flow.

- A *degraded restoration* occurs when a rerouted flow is degraded for a period (*T*) before it recovers its original reservation. Two forms of degraded restoration can occur: (i) a max-reserved mode flow operates at min-reserved mode and/or best effort mode and eventually recovers its original max-reserved mode service after some interval; (ii) a min-reserved mode flow operates at best effort mode and eventually recovers its original min-reserved mode service after some interval.
- A *permanent degradation* occurs when the rerouted flow never recovers its original reservation.

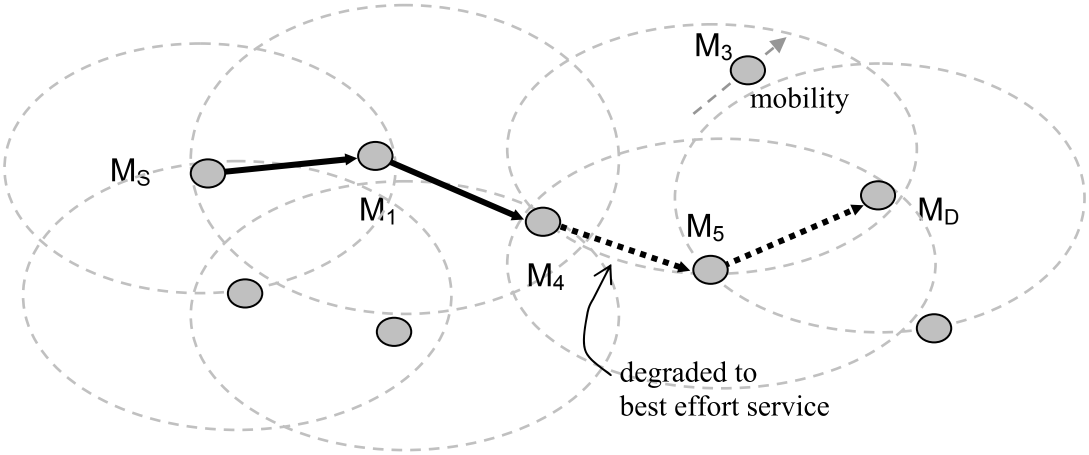

Figure 2-8: Rerouting and Degradation Illustration

Figure 2-8 illustrates the topology changes that occur after rerouting based on the initial topology shown in Figure 2-7. After rerouting link M_4-M_5 can only support best effort services. This type of restoration represents either a degraded restoration or a permanent degradation. In this scenario the destination node clears the partial reservation between mobile nodes M_S-M_4 by issuing a drop adaptation command to the source. The process of restoration can be immediate or delayed. Adaptation is application specific where the application can choose to respond to the network conditions and the delivered QOS.

2.5.2.5 Adaptation

The INSIGNIA QOS framework actively monitors network dynamics and adapts flows in response to observed changes based on user-supplied adaptation policy. Flow reception quality is monitored at the destination node and based on application-specific adaptation policy actions are taken to adapt flows under certain observed conditions. Action taken is conditional on what is programmed into the adaptation policy by the user. For example, an adaptation policy could be to maintain the service level under degraded conditions or scale down adaptive flows to their base QOS in response to degraded conditions; other policy aspects could be to always scale up adaptive flows whenever resources are available. The application is free to program its own adaptation policy, which is executed by INSIGNIA through interaction of the destination and source nodes.

INSIGNIA provides a set of adaptation levels that can be selected. Typically, an adaptive flow operates with both its base and enhanced components being transported with resource reservation. Scaling flows down depends on the adaptation policy selected. The flow can be scaled down to its base QOS delivering enhanced QOS packets in a best-effort mode, hence releasing any partial reservation that may exist. On the other hand, the destination can issue a drop command to the source to drop enhanced QOS packets (i.e., the source stops transmitting enhanced QOS packets). Further levels of scaling can force the base and enhanced QOS packets to be fully transported in best effort mode. In both cases, the time scale over which the adaptation actions occur is dependent on the application itself. These scaling actions could be instantaneous or based on a low pass filter operation [57].

During restoration of flow state, admission control and resource reservation are invoked. This can lead to changes in a flow's observed quality at the destination node both in terms of having to scale down flows in response to observed resource bottlenecks along the new path or scale up flows when additional resources are made available along the new path.

The INSIGNIA signaling system supports three adaptation commands that are sent from the destination host to the source using QOS reports:

- A *scale-down command* requests a source node to send its enhanced QOS packets as best effort or its enhanced QOS and base QOS as best effort.
- A *drop command* requests a source node to drop its enhanced QOS packets or enhanced and base QOS packets (where the term "drop" means the source node stop transmitting these packets).
- A *scale-up command* requests a source node to initiate a reservation for its base and/or enhanced service quality.

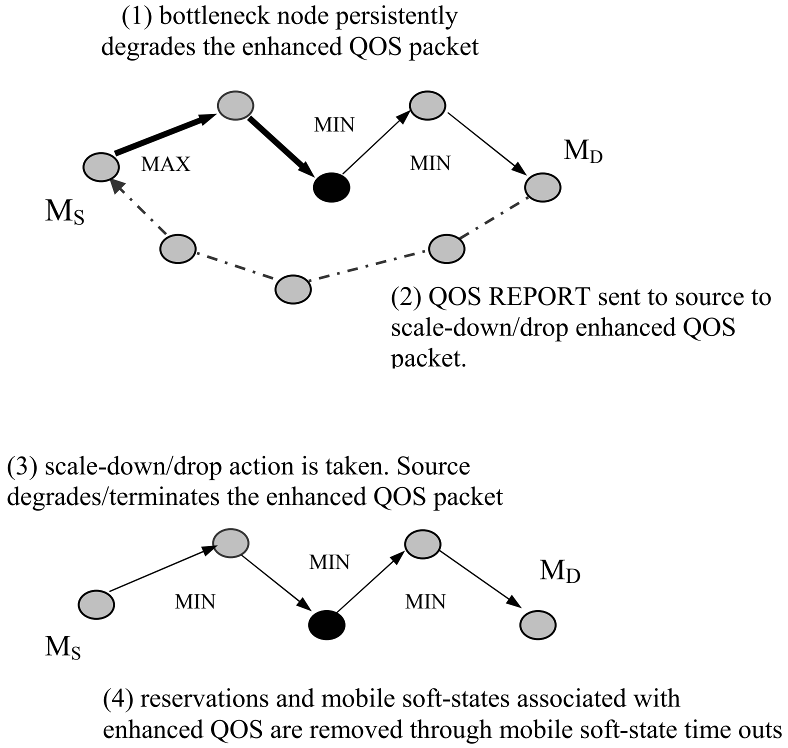
(1) bottleneck node persistently degrades the enhanced QOS packet

(2) QOS REPORT sent to source to scale-down/drop enhanced QOS packet.

(3) scale-down/drop action is taken. Source degrades/terminates the enhanced QOS packet

(4) reservations and mobile soft-states associated with enhanced QOS are removed through mobile soft-state time outs

Figure 2-9. Flow Adaptation

The scale down, drop, and scale up actions are driven by adaptation policy implemented at the destination, as illustrated in Figure 2-9. Note that preference is given to base over enhanced QOS components in the event reserved packets have to be degraded to the best effort mode at bottleneck nodes, as illustrated in the figure. The scale down command is issued when the degradation of enhanced QOS packets persists. This action forces source nodes to send the enhanced QOS packets as best effort packets, thereby effectively removing any partial reservations that may exist, as illustrated in Figure 2-9. A drop command is issued only when a destination node determines that degraded packets render insufficient quality. It is up to the applications to decide whether the reception of degraded packets is acceptable and take the appropriate action. An adaptation policy handler at the destination is free to issue scale down commands, or in the case of persistent degradation (possibly including best effort delivery of both the base and enhanced QOS components) to terminate the session.

Mobility results in the release of resources along old paths and session dynamic result in additional resources becoming available along existing paths when sessions terminate. These released resources help other source-destination pairs support higher levels of quality for their sessions assuming they share a common path with that of the released resources. In such a case, the signaling system sets the bandwidth indication in the packet's INSIGNIA IP option field to indicate to adaptation handlers (located at the receiver) that sufficient resources may be available to support the delivery of base and enhanced QOS. The signaling system uses the bandwidth indication field to inform the destination host of the availability of new network resources should they become available along an existing path. Bottleneck nodes set the bandwidth indicator to MIN when enhance QOS packets are scaled back in response to degraded conditions. Since each packet carries the max-min bandwidth requirements of each flow, bottleneck nodes can update a packet's bandwidth indicator in the event that resources become available to meet enhanced QOS needs. If all nodes along a path have resources to support enhanced QOS then the bandwidth indicator received at the destination will indicate MAX in the bandwidth indicator field. This does not imply that a reservation has been made or that a reservation could be made with a 100% assurance. Rather, it indicates to the source node that a reservation may be possible and that at the time the bandwidth indicator bit was set resources were available. To initiate the reservation for the enhanced QOS adaptation handlers send scale-up commands to their respective source nodes. In this sense the bandwidth indicator represents a good resource hint that additional service quality is possible. All messaging between source-destination pairs in support of scaling or dropping flow components is achieved using QOS reports.

2.6. Evaluation

In what follows, we present the evaluation of the INSIGNIA QOS framework through simulations, with emphasis on the performance of the signaling system. The goal of the simulations is to evaluate the suitability of the INSIGNIA to support adaptive flows in a mobile ad hoc network under various traffic, mobility, and channel conditions. In particular, we are interested in evaluating system-wide restoration and adaptation dynamics and the impact of soft-state mechanisms and mobility on end-to-end sessions.

2.6.1 Simulation Environment

The INSIGNIA simulator consists of 19 ad hoc nodes as illustrated in Figure 2-10. Each mobile node has a transmission range of 50 meters and shares a 2 Mbps air interface between neighboring mobile nodes within that transmission range. Time-varying wireless connectivity between nodes is modeled using 42 links. The mobility model is based on link failure and recovery characteristics defined in [61]; that is, connectivity is randomly removed and recovered with an arbitrary exponential distribution. Typically, mobile ad hoc networks do not have full connectivity between all mobile nodes at any given time due to the mobility behavior of mobile nodes and time-varying wireless link characteristics. With this in mind, maximum network connectivity is set at 85% such that 15% of the mobile nodes within their transmission ranges remain disconnected.

We discuss the implementation of our INSIGNIA QOS framework where the generic MANET routing protocol used is based on an implementation of the Temporally Ordered Routing Algorithm (TORA) [32].

The QOS architectural components implemented in our simulator include the following:

- The TORA [32] provided by the Naval Research Lab is used as a generic MANET routing protocol. The INSIGNIA framework is designed to "plug in" any MANET routing protocol.
- A packet scheduler, which based on a deficit round robin implementation [69].
- An admission controller, which is simply based on peak allocation of bandwidth.

For simulation purposes 10 adaptive flows with different bandwidth requirements ranging from 75-500 kbps are operational throughout the simulation. An arbitrary number of best effort flows are randomly generated to introduce different loading conditions distributed randomly throughout the network (i.e., in different parts of the networks) during the simulation. We also chose an arbitrary traffic pattern/load with average packet size of 2 Kbytes. Identical traffic/loads are used for all scenarios under investigation. The base QOS component of adaptive flows corresponds to 50-70% of an adaptive flow's bandwidth needs whereas enhanced QOS corresponds to 30-50%. For example, an adaptive flow of 300 kbps operating between nodes M_{14}-M_{13} (as illustrated in Figure 2-10) has 150 kbps for both its base and enhanced QOS such that minimum and maximum requirement is set to 150 kbps and 300 kbps, respectively.

The mobility model used throughout the simulations supports three different rates of mobility. Moderate mobility represents slow vehicular mobility ranging from 9-18 km/hr. Mobility conditions slower than moderate mobility is defined as slow mobility (i.e., speed less that 9km/hr) while rates faster than moderate mobility models are categorized as fast mobility (i.e., speed exceeding 18km/hr). We inherit the mobility model that was used in the TORA simulation [67]. In the simulation, we adopted a simple model for mobility pattern [67] that abstracts the mobility and wireless link characteristics into link failure and link recovery characteristics. A shortcoming of this approach is that mobile nodes have a fixed set of neighboring mobile nodes limiting the set of possible neighbors to communicate with. Therefore, the relative -and not absolute - mobility of the nodes is modeled. For the purpose of evaluating our framework, we measure per-session and aggregate network conditions for a number of experiments that analyze flow restoration, flow adaptation, soft-state management and host/router mobility. We observe throughput, delays, out-of-order sequence packets, lost packets, percentage of delivered degraded packets for the different mobility rates, and systems wide configuration (e.g., changing soft-state timers). We are particularly interested in percentage of reserved and degraded packets delivered to at all the receivers. This metric represents the ability of our framework to deliver assurance in mobile ad hoc networks. We also observe the number of rerouting, degradation, restoration, and adaptation events that took place during the course of each experiment as a measure of the dynamics of the system under evaluation

2.6.2 Restoration Analysis

In the following experiment we investigate the impact of rerouting and restoration on adaptive flows. Since rerouting of flows requires admission control, resource allocation, state creation, and removal of old state we track the rerouting and restorations events and any degradation that takes place. Typically, adaptive flows experience continuous rerouting during the session holding time. This is certainly the case with flows that represent continuous audio and video flows but not necessarily the case for microflows. These flows may be rerouted over new paths that have insufficient resources to maintain the required QOS. A key challenge for restoration is the speed at which flows can be restored. This is dependent on the speed at which new routes can be computed by the routing protocol if no alternative routes are cached and the speed at which the signaling system can restore reservations. The speed at which old reservations are removed is a direct function of the soft-state timer. The mobility rate impacts the number of restorations observed in the system and therefore the QOS delivered by the INSIGNIA QOS framework. As the rate of mobility increases (e.g., from moderate to fast), restoration algorithms need to be scalable and highly responsive to such dynamics in order to maintain end-to-end QOS.

In Section 2.5.2.4 we identified three types of restoration supported by the INSIGNIA model: immediate restoration, degraded restoration, and permanent degradation. Figures 2-10(a) and 2-10(b) illustrate the number of restorations and degradations that are associated with three randomly selected adaptive flows in our simulation. Due to the lack of resources at mobile node M_6, only flow M_{14}-M_{13} (i.e., the flow that traverses nodes M_{14}-M_{13}) is transported in max-reserved mode, while flows M_{16}-M_{18} and M_{15}-M_7 are transported in min-reserved mode. As a consequence, only the base QOS packets of flow M_{16}-M_{18} and flow M_{15}-M_7 are delivered as reserved mode packets, while enhanced QOS packers are transported as degraded best effort packets. As illustrated in Figure 2-10(b), flow M_{16}-M_{18} transported in min-reserved mode regains its max-reserved service through the rerouting of flow M_{15}-M_7. Rerouting of flow M_{15}-M_7 causes resources (i.e., 200 kbps) to be released by mobile soft-state management. Consequently, this action allows mobile router M_6 to restore the reservation requirement for the enhanced QOS of flow M_{16}-M_{18} which requires 80 kbps. The rerouting of flow M_{15}-M_7 finds sufficient resource availability on the new path (i.e., M_{15}-M_8-M_{11}-M_{18}-M_{14}-M_{10}-M_7), restoring its enhanced QOS.

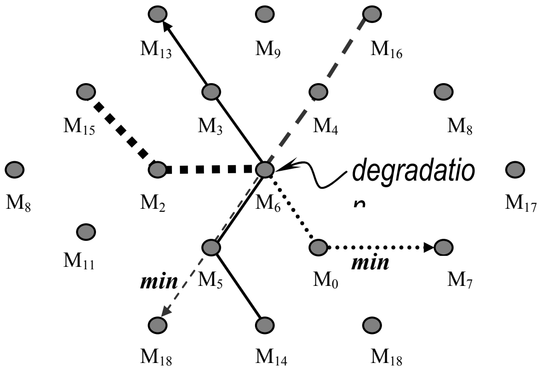

Figure 2-10(a): Degradation Due to Lack of Resources

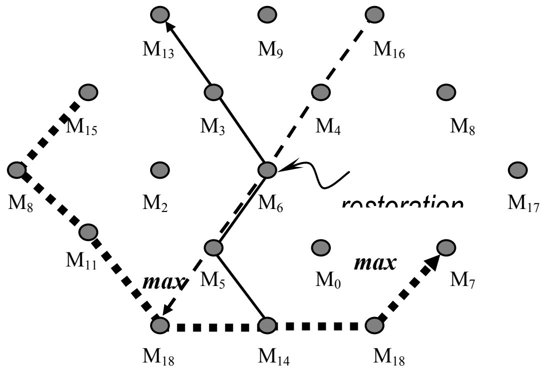

Figure 2-10(b): Restorations Through Rerouting of a Flow

Figures 2-11(a) and 2-11(b) illustrate immediate and degraded restorations observed under various mobility conditions. As indicated in the figures an increase in network dynamics increases the number of observed immediate and degraded restorations. The network experiences a total of 38 (61%) immediate restorations and 24 (39%) degraded restorations in the course of the simulation for a mobility rate of 3.6 km/hr, as illustrated in Figure 2-11(a). As mobility condition increases, the ratio between immediate restoration and degraded restoration changes. More immediate restorations are observed in comparison to degraded restorations for slow and moderate mobility conditions, as illustrated in Figure 2-11(b). However, when mobility conditions exceed 45 km/hr, degraded restoration becomes dominant as illustrated in Figure 2-11(b). The connectivity between mobile nodes becomes problematic as the mobility of nodes increases causing the network topology to rapidly change. Consequently, the number of available routes between source and destination nodes diminishes and the contention for network resources increases. This phenomenon introduces service fluctuations and degradation. Figure 2-11 illustrates the different types of restoration discussed in Section 2.5.2.4. Adaptive flows experience frequent re-routing with increased mobility causing a rise in the number of degraded restorations observed.

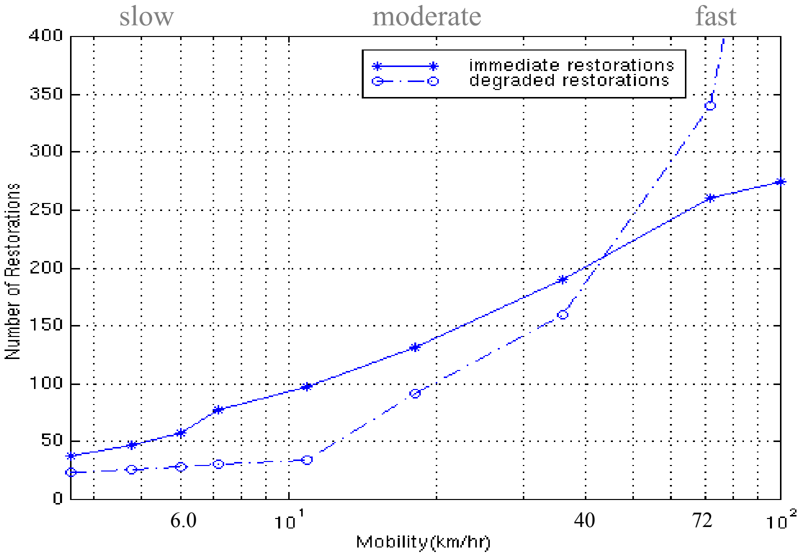

Figure 2-11(a): Number of Restorations

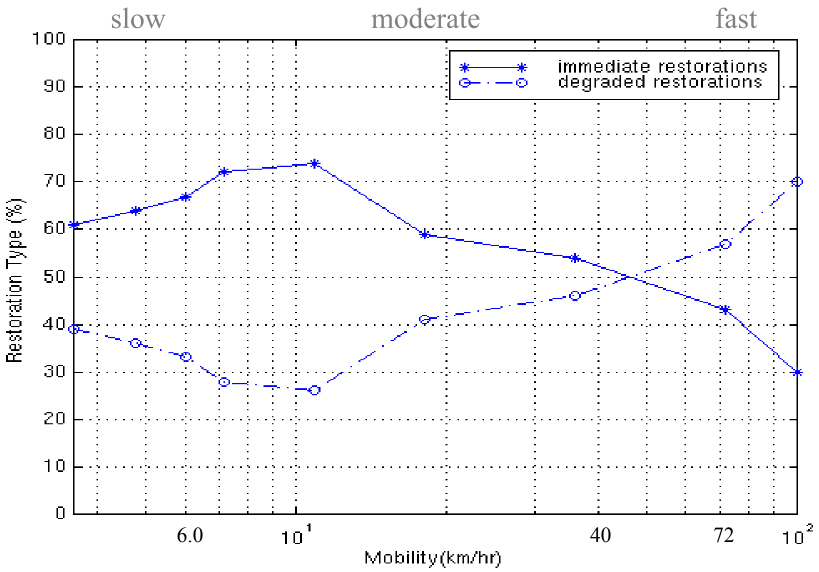

Figure 2-11(b): Percentage of Restorations

The INSIGNIA framework adopts a simple admission control test that does not favor rerouted flow over existing flows. A rerouted flow is denied restoration along a new route when insufficient resources are available to meet its minimum bandwidth requirements. This approach minimizes any service disruptions to existing flows, preventing a wave of service fluctuation to propagate throughout the network. When a mobile host loses its connectivity to neighboring nodes due to mobility, reservations along the old path are automatically removed. In the case of degraded restoration or permanent degradation, flows are degraded to min-reserved mode or best effort mode because of the lack of resources to restore the flows during rerouting. We observed that max-reserved adaptive flows are more likely to be degraded to best effort service than are min-reserved mode adaptive service. This is mainly due to the admission control policy adopted and semantics of base QOS and enhanced QOS components of flows where the base QOS of a typical adaptive flow consists of 50-70% of the overall bandwidth needs. The admission controller will attempt to support the base and enhanced bandwidth needs of flows. This leads to a situation where most mobile nodes mainly support max-reserved mode flows and a few min-reserved mode flows to fill the remaining unallocated bandwidth. This leads to the blocking of max-reserved flows and due to this behavior the vast majority of degraded flows are max-reserved to best effort. Therefore, degraded restorations of best effort to min-reserved (meaning that the min-reserved flow is degraded to best effort before being restored to min-reserved) only occur when the rerouted adaptive flows encounter resources to support only min-reserved service. We observed that degraded restoration for best effort to max-reserved (meaning that the max-reserved flow is degraded to min-reserved and/or best effort before being restored to max-reserved) is the most dominant degraded restoration type observed, as shown in Figure 2-12. This is because rerouted flows are more likely to be accepted or denied rather than degraded to min-reserved flows under slow and moderate mobility conditions. However, we observe that when the mobility exceeds 72 km/hr that best effort to min-reserved degraded restoration becomes the dominant type, as shown in Figure 2-12. In the case of high mobility, only a limited number of routes exist to route flows, causing service degradation. The rapid fluctuations in the monitored QOS cause the adaptation processes at the destination to request that the degraded flows be scaled down to their min-reserved mode. In this instance, the best effort to min-reserved restoration becomes the dominant type, as shown in Figure 2-12.

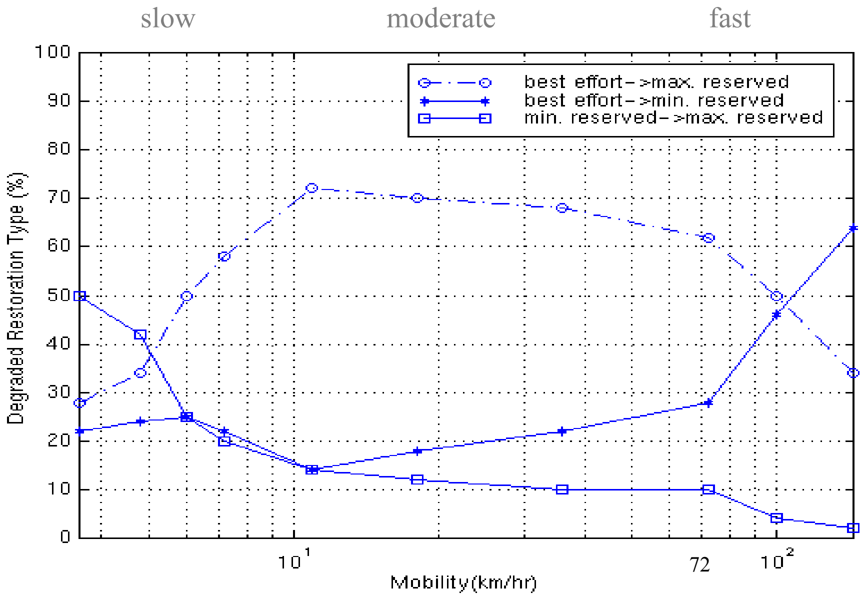

Figure 2-12: Degraded Restorations Types

Increased mobility forces mobiles hosts to adapt flows to their min-reserved modes and prevents adaptive flows from scaling back up due to the fast time scale dynamics and rerouting observed. When the mobility exceeded 72 km/hr, all adaptive flows are scaled down to their min-reserved service 90 seconds into the trace. Only two scale up adaptations actions were observed during the complete trace. The number of best effort to max-reserved and min- reserved to max-reserved degraded restoration types decrease as mobility is increased beyond 72 km/hr, as shown in Figure 2-12. The best effort to min-reserved degraded restoration continues to increase, implying that most of flows scale down to their minimum requirements and operate at the min-reserved mode.

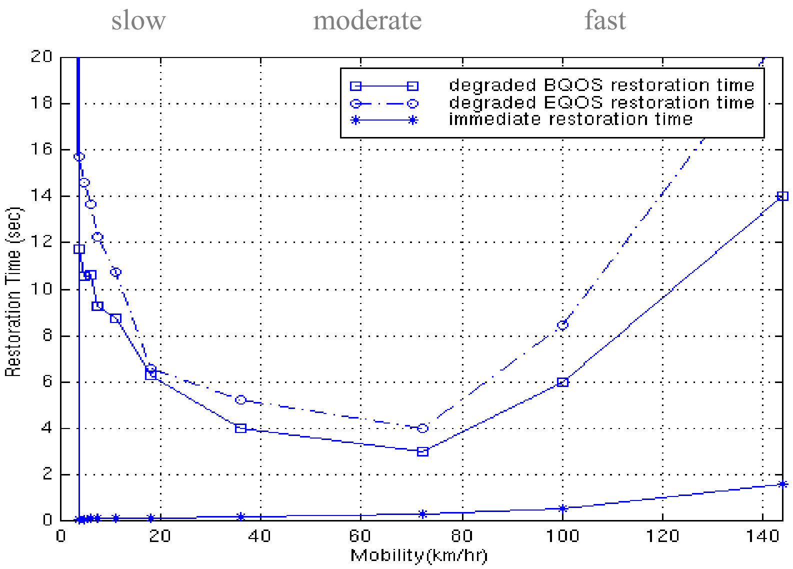

Figure 2-13: Time Spent for Immediate Restorations and Degraded Restorations

Figure 2-13 shows the restoration times across the complete mobility range. The base QOS restoration time corresponds to the time taken to regain the min-reserved service for a flow that has been temporarily degraded to a best effort mode service. The enhanced QOS restoration time corresponds to the time taken for the max-reserved service to restore from the best effort service or from min-reserved service. We observe that the average required restoration time for immediate restoration is relatively constant at 0.2 ~ 0.9 seconds under all mobility conditions. We observe that immediate restoration only require an interval of two consecutive packets to restore the reservation. However, mobility conditions impact the average degraded restoration times, unlike the immediate restorations, as shown in Figure 2-13.

2.6.3 Adaptation Analysis

The adaptation process operates on an end-to-end basis and is driven by the observed service quality and adaptation policy of the destination node. This is in contrast to restoration, which operates on the re-routing time scale. Typically, adaptation operates over longer time scales associated with end-to-end applications and their adaptation strategies. Monitoring modules residing at destination nodes actively measure the delivered service quality. As discussed in Section 2.5.2.5, destination nodes can issue adaptation commands to source nodes using QOS reports to scale down, drop and scale up flows. For example, when the degradation of enhanced QOS packets persists beyond an acceptable period, the destination can issue a scale down adaptation command to the source node removing any partial reservations that may exist between the source host and the bottleneck host. The INSIGNIA system is also capable of scaling up flows (e.g., from a min-reserved to a max-reserved service). The bandwidth indicator plays a central role in the adaptation process, as discussed in Section 2.5.2.5.

To observe the dynamics associated with the adaptation process, two adaptive flows are arbitrarily chosen and their associated throughputs measured (at their destination nodes) over the course of the simulation. The simulation results reflect moderate mobility conditions of 11 km/hr. Moderate mobility conditions were chosen because slow mobility lacks network dynamics and fast mobility rarely experiences end system-initiated adaptation due to the rapid fluctuations in resource availability.

The impact of the adaptation process, degradation, and restoration on flows M_{15}-M_7 and M_{16}-M_{18} from the previous example is shown Figure 2-14. As shown in the trace, flow M_{16}-M_{18} is affected by network dynamics at 17 seconds into the trace. The mobility of the network forces flows to be rerouted and, due to lack of resources along the new path, causes flow M_{16}-M_{18} to degrade to the min-reserved service, as indicated by (1) in Figure 2-14. The degradation of flow M_{16}-M_{18} enhanced QOS packets is restored at (2) in Figure 2-14. Degradation of the base QOS at point (3) is observed at 160 seconds and it is preceded by degradation of enhanced QOS packets at 145 sec into the trace. Due to persistent service disruption the destination node (M_{18}) triggers the source node (M_{16}) to scale down the flow at 151 sec into the trace. The decision to scale down the flow is controlled by an adaptation handler. The source responds by transmitting the enhanced QOS packets as best effort packets. The reservations associated with the enhanced QOS packers is de-allocated by soft-state management operating at intermediate routing nodes along the path, allowing other adaptive flows to scale up. Scaling up can be observed at t=172 seconds into the trace when the destination node (M_{18}) detects consistent resource availability through monitoring the bandwidth indicator. Flow M_{16}-M_{18} restores its max-reserved mode service while flow M_{15}-M_7 first experiences degradation, scaling down, and then scaling up. The degradation of flow M_{15}-M_7 enhanced QOS packets degraded at t = 92 seconds is restored (2’) to max-reserved mode service at t = 98 seconds into the trace. However, further network dynamics force the degradation of the enhanced QOS packet at t = 100 seconds into the simulation.

Adaptation policy is application specific in the sense that some flows prefer to instantly scale up when resources become available while others prefer not to follow instantaneous changes but trends in resource availability. The scaling policy can be based on simple algorithms, for example, a simple state machine that scale flows down or up based on a certain number of degraded packets or packets indicating that additional resources are available, respectively. More sophisticated algorithms could follow statistical observations about network dynamics using low pass filters.

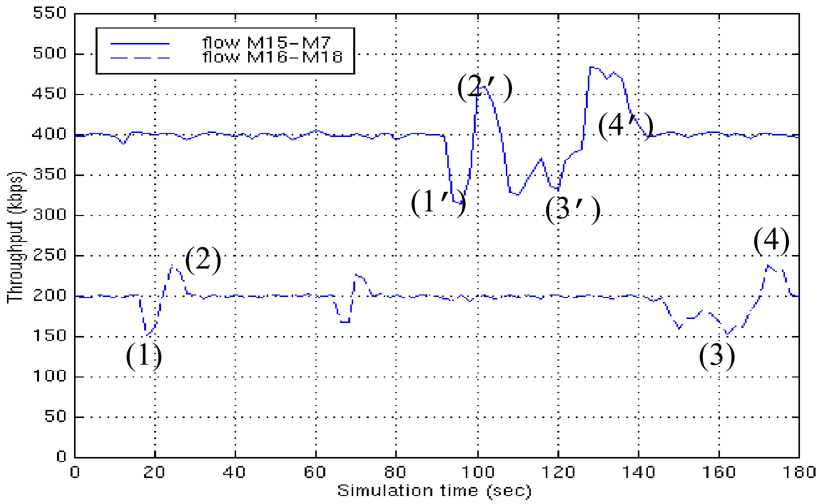

Figure 2-14: Trace of Adaptive INSIGNIA Flows

The rate of mobility has a large impact on the observed adaptation dynamics. Fewer instances of adaptation are observed given the same adaptation policy for slow mobility over moderate mobility. For mobility of 3.6 km/hr we observe two scale-up actions and one scale down action, whereas at 18 km/hr we observe seven scale-up and four scale-down actions. As mobility increases beyond the moderate rate we observe more fluctuation in delivered service quality where scaling down flows to a min-reserved service becomes common. As the mobility speed increases to fast we observe few scaling up actions due to the fast dynamics of the network. Few destinations observe stable conditions to issue a scaling up command to their peer source nodes. For example, at 72 km/hr we observe that only two scaling up actions are recorded, with all adaptive flows being forced to scale down to their min-reserved mode during the course of the simulation.

2.6.4 Soft-State Analysis

Soft-state resource management is used to maintain reservations. The duration of soft-state timer has a major impact on the utilization of the network. Figure 2-15 shows the impact of soft-state times on network performance in terms of the number of reserved mode packets delivered. Reception of a reserved mode packet (with the service mode set to RES, as discussed in Section 2.5.1.1) at the destination indicates that the packet is delivered with max-reserved or min-reserved assurance. Reception of a packet degraded implies that the packet has been delivered without such guarantees. Therefore the percentages of reserved and degraded packets received by destination nodes as a whole indicate the degree of service assurance that an INSIGNIA network can support for different values of soft-state timers.

In what follows, we discuss the impact of soft-state timers on network performance. We set the mobile soft-state timer value in the range of 0.01 to 30 seconds and observe the corresponding system performance. For each experiment we set the same timer value at each node. As shown in Figure 2-15, the mobile soft-state timer value has an impact on the overall network performance. The ability to support adaptive services decreases as the soft-state timer value increases. The percentage of delivered reserved packets decreases as mobile soft-state timer increases. The percentage of degraded packets increases as the soft-state timer value increases, as shown in Figure 2-15. Worst performance is observed when the soft-state timer value is set to 30 seconds. In contrast, the best performance is observed when mobile soft-state timer is set to 2 seconds, as shown in Figure 2-15. We observed that 69% of the packets are delivered as reserved packets and 31% as best effort packets when the soft-state timer is set at 30 seconds. Support for QOS substantially improves with 88% of reserved packets being delivered to the receivers with a soft-state timer value of 2 seconds. Large timeout values tend to lead to underutilization of the network because resources are "locked up" where resources remain allocated long after flows have been rerouted. New flows are unable to use these dormant resources, resulting in the overall degradation of the network due to "resource lockup".

As the value of the soft-state timer gets smaller fewer resource lockups are observed and utilization increases. However, when the timer is set to a value smaller than 2 seconds the network experiences what we describe as "false restoration". This occurs when a reservation is prematurely removed because of a small soft-state timer. However, this is a false state because the session holding time is still active and the source node keeps sending packets. In this case, the reservation is removed because of a timeout and then immediately reinstated when the next reserved packet arrives.

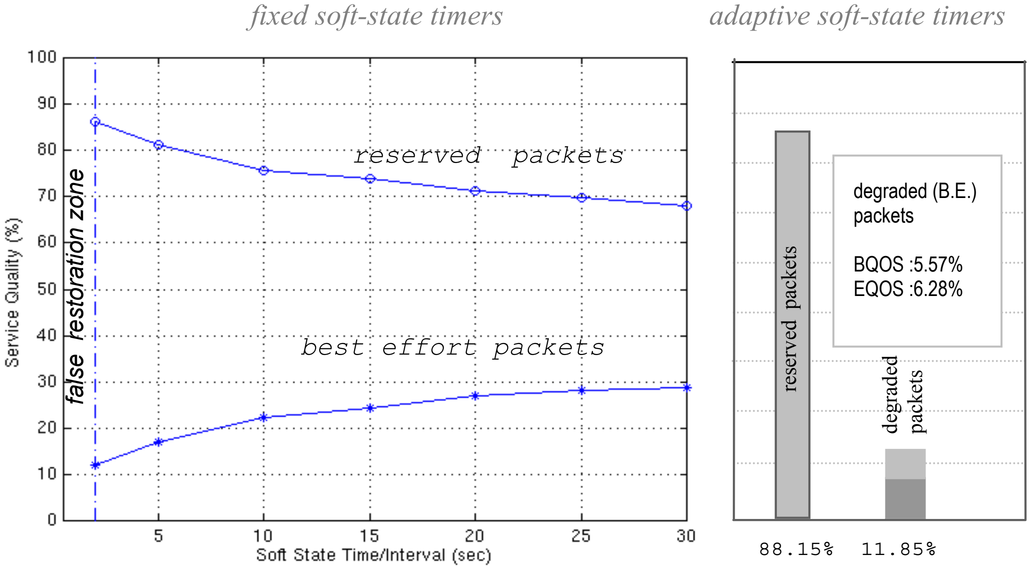

fixed soft-state timers

adaptive soft-state timers

Figure 2-15: Soft-state Timers and Network Performance

False restorations occur when the timeout value is smaller than the inter-arrival time between two consecutive packets associated with a flow. With a soft-state timer of 0.04 sec, for example, all the adaptive flows experienced numerous false restorations. Mobile routers often deallocate and reallocate resources without the involvement of any network dynamics due to mobility. In the worst case, every packet can experience a false restoration. Such events not only increase the processing costs of state creation and removal, and resource allocation and deallocation, but also falsely reflect the resource utilization and availability of the system. When the network experiences numerous false restorations, rerouted flows often find nodes with few resources allocated on the new path. This phenomenon causes flows to always gain max-reserved mode resources with mobile nodes accepting the request for resources well beyond their actual capacity. This results in reserved packets experiencing indefinite delays at intermediate nodes even though resource assurances are provided by admission controller, resulting in wide scale packet loses and service degradation. Figure 2-15 shows a “false restoration region” where there is little distinction between reserved and best effort operational modes and where reservations are typically always granted. Adaptation and restoration algorithms can fail under false restoration conditions due to perception of unlimited resource availability. Setting a suitable soft-state timer value is therefore essential to preventing both false restoration and resource lockup in our framework.

Each data packet associated with a reserved flow is used to refresh soft-state reservations. We observe that different adaptive flows have different data rates, and thus a fixed timeout value is too limiting. For example, one value may be fine for some set of flows but cause false restorations or resource lockup for others. Clearly there needs to be a methodology for determining the value of the soft-state timer. The issue of false restoration and resource lockup can only be resolved by adjusting the timeout value based on the observed flow dynamics. The timeout should be based on the effective data rate of each flow. More specifically, the soft-state timer should be based on the measured packet inter-arrival rate of adaptive flows. The signaling system measures packet inter-arrivals and jitter at each mobile node for each flow, adjusting the soft-state timeout accordingly. In the experimental system we implemented an *adaptive soft-state timer* that is initially set to 4 seconds, representing an initial safety factor. This allows mobile nodes to set their soft-state timers according to their effective data rate, allowing the timeout to adjust to network dynamics and the variation in the inter-arrival rates of individual flows traversing nodes. The implementation of an adaptive soft-state timeout effectively removes resource lockups and false-restorations, as shown in Figure 2-15. We observe that when an adaptive soft-state timer scheme is used 88% of flows are delivered as reserved packets and 11% as degraded packets. Adaptive soft-state timers greatly reduce resource lockup and false restoration conditions, allowing the network to support better service assurances through the delivery of more reserved packets and fewer degraded packets at destination nodes.

2.6.5 Mobility Analysis

To evaluate the impact of mobility on the INSIGNIA QOS framework, we conduct a set of experiments operating under identical traffic patterns/load conditions and various mobility conditions ranging from 0 km/hr to 72 km/hr. Figure 2-16 illustrates the impact of mobility on the delivered service quality. When there is no host mobility, results closely approximates a fixed network infrastructure where admitted flows receive stable QOS assurances. One anomaly is observed, however. Six adaptive flows failed to be granted reservations due to lack of network resources at intermediate nodes. As consequence only 49% of the packets are delivered as reserved packets and 51% as best effort packets. This anomaly is a product of the routing protocol, which provides a non-QOS routing solution. Adaptive flows are routed to bottleneck nodes resulting in the failure of admission control due to the lack of resources. This problem could be resolved by designing a signaling system that takes alternative routes in the case that admission control fails along a selected path.

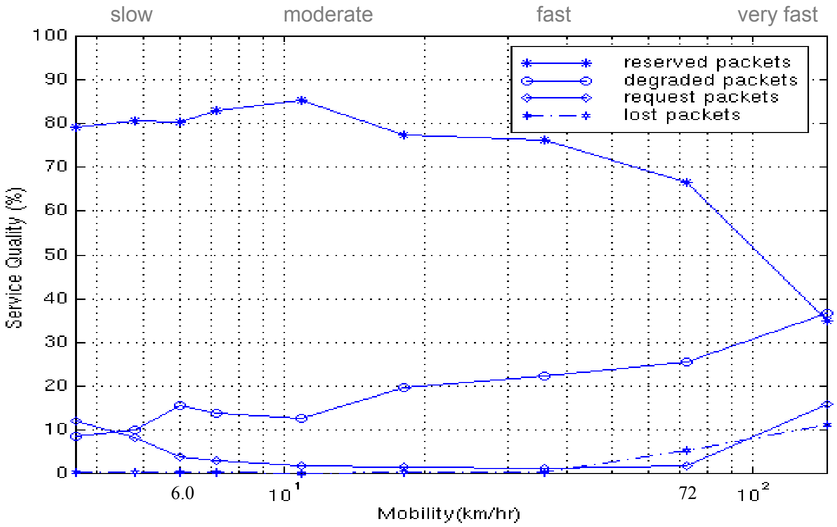

Figure 2-16: Mobility and Network Performance

With the introduction of mobility into the network, the performance improves (i.e., more reserved packets are delivered) as illustrated in Figure 2-16. Mobility-induced rerouting allows request packets to traverse alternative paths, increasing the probability of finding a route with sufficient resource availability to admitted flows as reserved mode packets. Figure 2-16 shows that INSIGNIA supports relatively constant QOS under slow and moderate mobility conditions between 3.6 and 18 km/hr. The optimal performance is observed when the average network mobility is approximately 11 km/hr. This results in the delivery of 86% of reserved packets. The in-band nature of INSIGNIA allows the system to cope with fast network dynamics in a responsive manner. In an ideal case, INSIGNIA requires only a single packet reception to set up and restore (i.e., immediate restorations) reservation for the new or rerouted flows, respectively. INSIGNIA supports the delivery of 66% reserved packets even when mobiles are moving at 72 km/hr as shown in Figure 2-16. This is a very encouraging result.

Note that the service provided in a mobile ad hoc network has a memoryless property such that adaptive flows require new admission tests along the new path when rerouting occurs. This implies that an increase in mobility may cause fluctuations in perceived service quality. At 72 km/hr all flows are scaled-down to min-reserved packets after 90 sec into the simulation due to the fluctuations in delivered quality. At this speed only two flows are capable of regaining their max-reserved service. When mobility conditions exceed 72 km/hr, support for QOS breaks down rapidly as indicated in Figure 2-16. The mobility characteristics overload the system and service assurance for adaptive flows diminishes. In fact, when mobility exceeds 90 km/hr, we observe that flows M_{12}-M_{11}, M_3-M_7 and M_5-M_{12} are transported as best effort packets for more than 70 seconds because they failed to accomplish their end-to-end flow set up due to persistent loss of RES packets and QOS reports. This phenomenon corresponds to the abrupt loss of reserved packets and degraded packets.

An increase in out-of-sequence packet is also observed at higher speeds, possibly causing service disruption at the receiver. Figure 2-17 shows the number of out-of-sequence packets under various mobility conditions. The number of out-of-sequence packets generally increases as mobility increases. The number of delivered out-of-sequence packets is impacted by different propagation delay characteristics of reserved and best effort packets associated with the same end-to-end flow. Figure 2-17 also shows the number of lost packets observed under different mobility conditions. Packets that are delayed for more than 15 seconds are discarded at intermediate nodes and considered lost. Figure 2-18 shows the delay characteristics of packets under various mobility conditions. When mobility increases, the connectivity between nodes becomes problematic. Such network dynamics trigger frequent routing updates and decreased connectivity. Thus, the number of available routes between nodes decreases as mobility increases. Degraded packets queue up at intermediate nodes experiencing long delays. However, the reserved packets are less sensitive to these delays, as indicated in Figure 2-18, with all reserved packets being delivered within a period of 40 milliseconds.

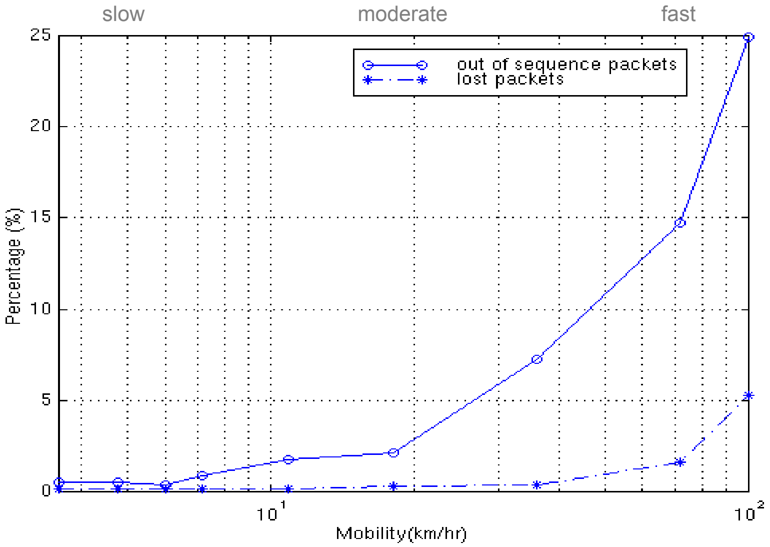

Figure 2-17: Impact of Mobility on Out of Order Delivery and Packet Loss

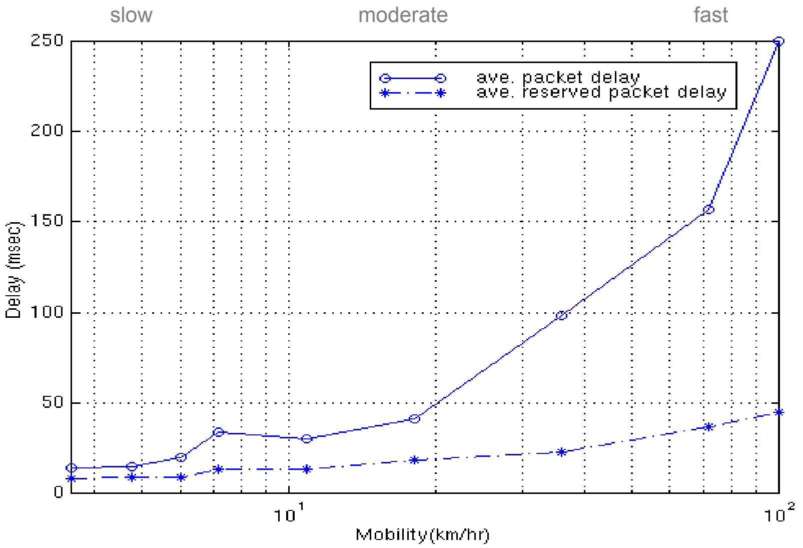

Figure 2-18: Average Packet Delays

2.7. Conclusion

In this chapter we have presented the design, implementation, and evaluation of the INSIGNIA QOS framework that supports the delivery of adaptive services in mobile ad hoc networks. A key contribution of our framework is the INSIGNIA signaling system, an in-band signaling system that supports fast reservation, restoration, and adaptation algorithms. The signaling system is designed to be lightweight and highly responsive to changes in network topology, node connectivity, and end-to-end quality of service conditions. We have evaluated our QOS signaling framework paying particular attention to the performance of the signaling system.

The approach discussed in this chapter looks promising in terms of performance results presented. Our simulation results show the benefit of our framework under diverse mobility, traffic, and channel conditions. The use of in-band signaling and soft-state resource management proved to be very efficient, robust, and scalable. Our results highlighted a number of anomalies that emerged during the evaluation phase. However, the use of adaptive soft-state timer seemed to resolve many of these issues (e.g., false restorations and resource lockups).

Based on the adaptive QOS framework introduced in this chapter, we present a detailed evaluation of the INSIGNIA signaling system in Chapter 3. Specifically, we investigate how well INSIGNIA performs with a number of MANET routing protocols and supports the adaptive QOS for TCP/UDP flows in diverse networking conditions.

Chapter 3

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

3.1 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, the Army Research Laboratory Federated Laboratory and the DARPA Global Mobile Information Systems (GloMo) program [59]) 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 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 mobile ad hoc networks 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 on-going 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 routing (AODV) [30], Dynamic Source Routing (DSR) [31] and Temporally Ordered Routing (TORA) [32] among others [56]. 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 [45] [51] [52] [61] reporting on the performance of AODV, DSR and TORA in the context of best-effort networks.

The contribution of this chapter is as follows. Section 3.2 reviews the INSIGNIA signaling system, and Section 3.3 describes our *ns-2* [40] simulation environment used for the evaluation of the system. We evaluate the performance of INSIGNIA to seamlessly interoperate with AODV [30], DSR [31], and TORA [32] showing that 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 in Sections 3.4

and 3.5, respectively. Performance of the restoration algorithm relies on the speed at which routing protocols can re-compute 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. Section 3.6 discusses our results and presents some concluding remarks.

3.2 INSIGNIA Overview

The INSIGNIA signaling systems provides support for adaptive reservation-based services in mobile ad hoc networks. 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 are encoded using the IP option field in datagrams. This in-band information is 'snooped' as data packets traverse intermediate nodes/routers and is used to maintain 'soft-state' reservations in support of flows/microflows.

3.2.1 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 packets 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 Figure 2-3 (see Chapter 2). 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 mobile ad hoc networks is the ability of the MAC layer to deliver service quality. INSIGNIA is an end-to-end IP-based reservation mechanism that is 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 Section 3.3, we outline a modification to the IEEE 802.11 [50] MAC distributed control function (DCF) that offers a simple set of differentiated services that INSIGNIA is build on.

A source node continues to sends packets with the reservation request bit set until the destination node completes the reservation set-up phase by informing the source node of 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, this 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 3-1(a) 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 M_1 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 Figure 3-1(a).

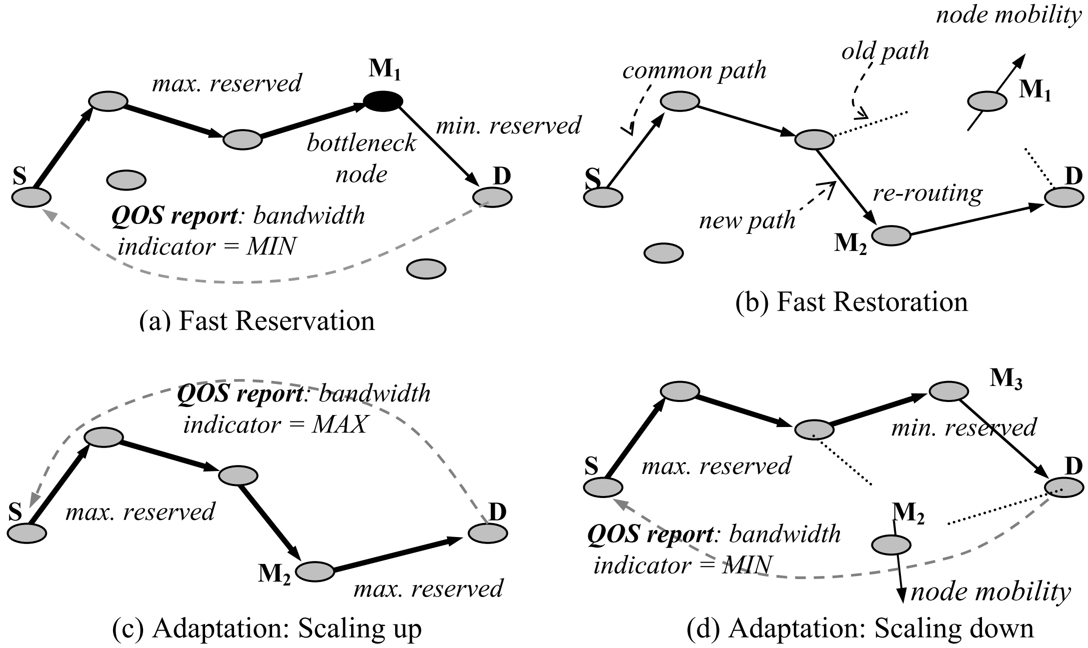

(a) Fast Reservation

(b) Fast Restoration

(c) Adaptation: Scaling up

(d) Adaptation: Scaling down

Figure 3-1: Examples of INSIGNIA Operations

INSIGNIA is designed to operate over unidirectional and bi-directional 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 node 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 are active at all intermediate routers.

3.2.2 Fast Restoration

Reservation-based flows are often re-routed within the lifetime of on-going sessions due to node mobility, as illustrated in Figure 3-1(b). In such cases, INSIGNIA performs fast restoration. The goal of restoration is to re-establish reservations as quickly and efficiently as possible. Re-routing active flows involves the MANET routing protocol (to determine new routes), admission control and resources 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 re-routed 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 re-routed to a new node where resources are unavailable, the flow is degraded to a 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 de-allocated. A reservation may be restored if resources are freed up at a bottleneck node or further re-routing 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 3-1(b) illustrates a fast restoration scenario where an intermediate node M_1 moves out of radio contact and a reservation-based flow is re-routed through the 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 Figure 3-1(b). 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 re-routing or session termination) then soft-state timers expire and resources are de-allocated. 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 Figure 3-1(b) timeout automatically. In this case, explicit signaling would not work because M_1 is out of radio contact form 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 [68] we report that an adaptive soft-state timer approach resolves a number of pathologies found in reservation-based mobile ad hoc networks such as 'false restoration' and 'resource lock-up' which limit performance.

3.2.3 End-to-End Adaptation

The INSIGNIA system supports on-going 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. Other 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 the 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 Figure 3-1(c). Node mobility or session dynamics cause a flow routed via M_2 to be scaled up from a minimum to maximum reserved service. The destination node (D) notes that the bandwidth indicator bit changes from a MIN to MAX value. This indicates that the current path could support higher levels of service. This indication is a really hint from the network (and not an 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. In this example, we show 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 Figure 3-1(d) a flow receiving maximum service is re-routed due to the mobility of node M_2. The new path through node M_3 has insufficient resources to support the maximum reserved 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, etc.). 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 that the application considers appropriate.

In essence, INSIGNIA provides a simple API to the network to implement sophisticated adaptation policies at the edge (i.e., source/destination) in a scalable manner.

3.3 Simulation Environment

In what follows, we discuss our simulation environment used to assess the performance of UDP and TCP over INSIGNIA-enabled mobile ad hoc networks. The full INSIGNIA code suite and test scripts used for the evaluation of the system are freely available on the Web [71]. The simulation environment uses the NS-2 [40] simulator and its wireless extensions developed by Monarch Project [55]. In this chapter, we use the terms 'INSIGNIA system' and 'best-effort system' to refer to the AODV, DSR and TORA networks with and without INSIGNIA support, respectively. In Section 3.4, 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 meters and shares a 2 Mbps radio channel with its neighboring nodes. We use a random way-point mobility model [61] 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 discussed in this chapter represent per-mobile packet generation intervals (e.g., 0.1 represents 10 packets/sec per mobile host). The simulated network area has a rectangular shape of 1500 meters by 300 meters that minimize 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 [73] includes the signaling system and a number of framework mechanisms discussed in [68]. 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 [68] 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 [74], we describe a MAC layer based on modifications to the IEEE 802.11 distributed control function that provides simple differentiated service. The MAC ensures that not only packets sent by the mobile host itself are differentiated, but more importantly, 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 [74], is possible by appropriately adjusting the back off times through the contention window limits. Two classes of services 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 [74].

Twenty flows are active during the simulation and are started with staggered times. Six of these flows are arbitrarily selected and monitored for the duration of 300 seconds 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 Kbps to 1.39 Mbps. The network resources are partitioned a priori such that at most 800 Kbps is allocated for reservation-based flows with the remainder supporting best effort traffic. This 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 sec pause time, which represents no mobility, to 0 sec pause time, which represents continuous mobility with a maximum speed of 20 m/sec (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 toward UDP and TCP performance.

3.4 UDP Performance

Previous performance comparisons [30] [31] [32] of AODV, DSR and TORA in best effort networks have often favored lightly loaded networks with relative small packet sizes. As a result, measured performance often achieves over 90% in packet delivery fraction; that is, 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 mobile ad hoc networks, however. In this section, we evaluate the performances 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.

3.4.1 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 Figure 3-2. 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 sec) while the performance of the six monitored flows is observed. Identical simulations were conducted for AODV, DSR and TORA networks showing the operational transparency of INSIGNIA to work with these routing protocols and to observe the performance differences that exist among these different MANET routing protocols.

As illustrated in Figure 3-2, the best-effort system (represented by the dotted lines in the plots) achieves more than a 90% packet delivery fraction under lightly loaded network conditions. This is consistent with results found in the literature [51] [52] [61]. 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% for all MANET routing protocols when the cross traffic exceeds 716 Kbps representing a packet generation interval of 0.08 sec. In addition, less than 60% of the packets are delivered when the cross traffic increases to 1.14 Mbps representing a packet generation interval of 0.05 sec.

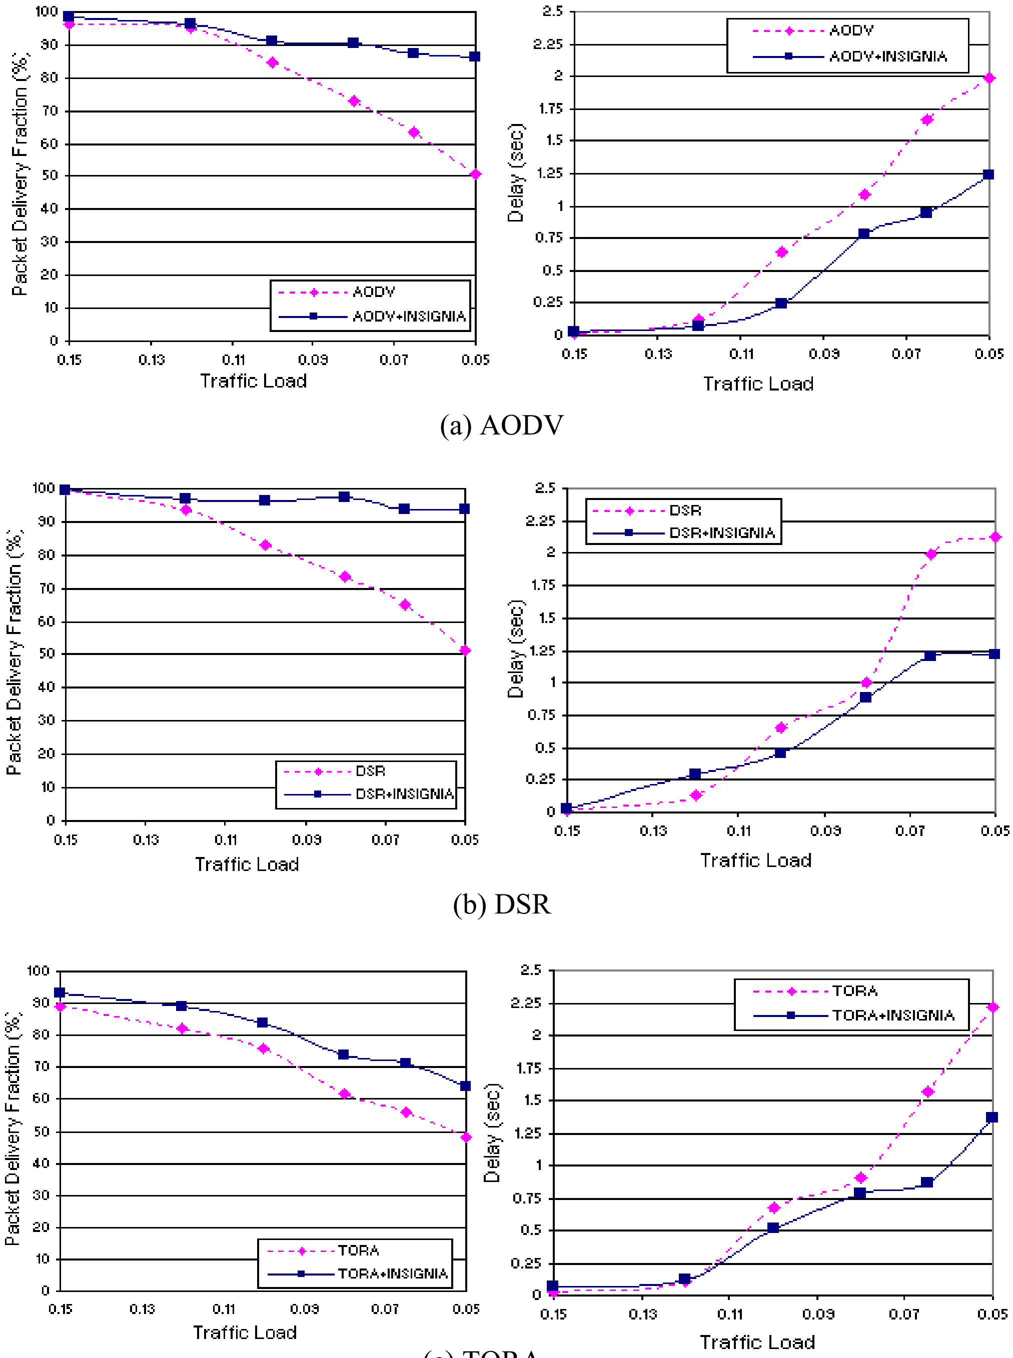

(a) AODV

(b) DSR

(c) TORA

Figure 3-2: Comparison of the Best Effort and INSIGNIA Systems under Increasing Network Load

Corresponding end-to-end delay measurements show a substantial 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 Figure 3-2. The performance improvements of the INSIGNIA system are shown in comparison to the best-effort system for each of the MANET routing protocol (viz. AODV, DSR and TORA) under study.

As shown in Figure 3-2, 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% when a cross traffic load of 573 Kbps (represented by packet generation interval of 0.10 sec) is introduced into the best-effort system. As cross traffic load increases to 1.14 Mbps (represented by packet generation interval of 0.05 sec), only 60% of the packets are delivered. In contrast, the packet delivery fractions for reservation-based flows do not drop below 88% 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 Figure 3-2. We observe that under lightly loaded condition 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% improvement in the INSIGNIA system for all the MANET routing protocols, as shown in Figure 3-2. 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 under performs due to the number of signaling messages generated to create and maintain valid routes.

3.4.2 Impact of Mobility

The impact of node mobility on the performance of the best-effort system is shown in Figure 3-3. The simulation duration is set to 300 seconds with 20 flows active in the network. We use the same mobility metric (i.e., pause time) defined in [61] to align our simulation results. The maximum and minimum mobility conditions are represented by a pause time of 0 and 300 seconds, respectively. The effect of mobility is observed by gradually decreasing the pause time of mobile nodes with the traffic load fixed at 800 Kbps (i.e., 40 Kbps/flow). As shown in Figure 3-3, as mobility increases packet loss and the 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 burst after topology changes. Such conditions are very difficult to control and provision for in ad hoc mobile networks. This inevitably leads to degraded restoration of re-routed 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 toward service disruption and degradation at the destination. While many flows experience degraded service quality when mobility increases, some of 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 re-routed out of a congested portion of the network as mobility increases. This phenomenon is also observed in [52].

Figure 3-3 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 3-3 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% improvement in the packet delivery fraction for AODV and DSR under low mobility condition and more than 7% for the TORA protocol. As mobility increases, the benefits of INSIGNIA over the best-effort network narrows, as shown in Figure 3-3. 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 re-routing 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.

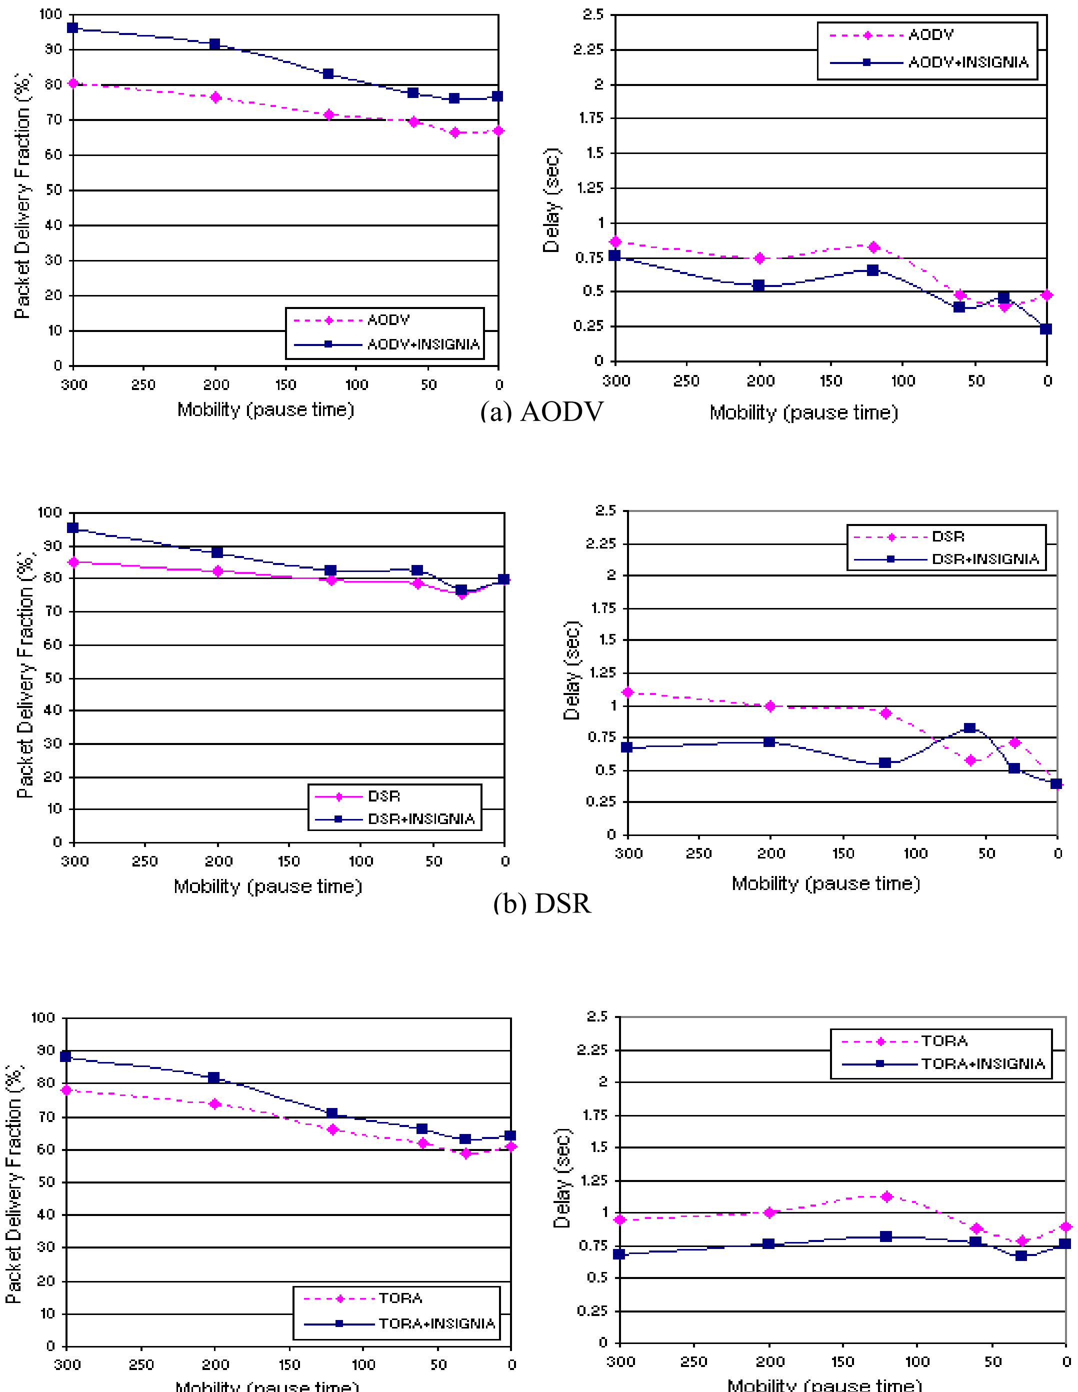

(a) AODV

(b) DSR

(c) TORA

Figure 3-3: Comparison of the Best Effort and INSIGNIA Systems under Increasing Node Mobility

The end-to-end delay measurements of the monitored flows in the INSIGNIA system also show improvement in comparison to the best-effort system, as shown in Figure 3-3. 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 % when the mobility is set at 200 sec pause time in contrast to 77 % in the best-effort system. Therefore, the average end-to-end delay measurement of 0.75 seconds in the best-effort system corresponds to the 80% packet delivery fraction while the average end-to-end delay measurement of 0.51 seconds in the INSIGNIA system corresponds to the 92% packet delivery fraction. The INSIGNIA system not only decreases the packet loss but also reduce the end-to-end delay.

Figure 3-4 and Figure 3-5 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 Kbps UDP/CBR flow operating under low to moderate mobility conditions (i.e., 120 sec pause time). The bandwidth requirement for the flow is defined by a minimum data rate of 22 Kbps. Figure 3-4 shows the throughput trace of the flow in the best-effort system and Figure 3-5 shows the throughput trace of a reservation-based flow in the INSIGNIA system. The monitored flow is rerouted 6 times during the simulation period and traverses 3 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 Kbps. In addition, 43%

of the transmitted packets are lost and 65% of the delivered data packets exceed 800 milliseconds end-to-end delay. In contrast, near constant rate throughput is observed for the same flow in the INSIGNIA system with 2% packet loss and only 9 % of delivered packets exceeding 800 milliseconds end-to-end delay.

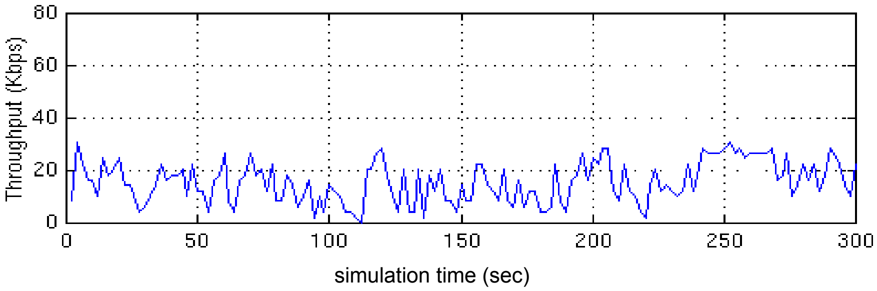

Figure 3-4: Trace of a Monitored Flow in a Best-Effort System

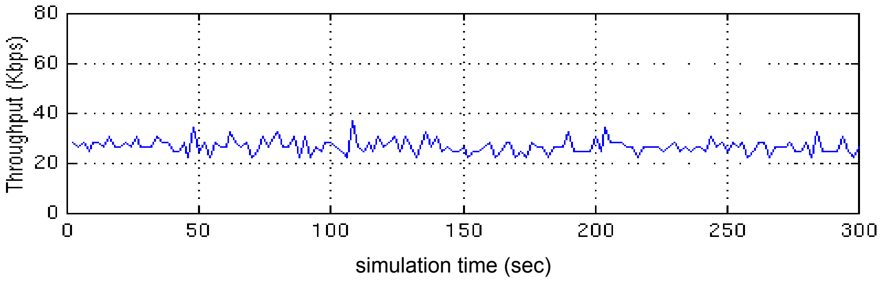

Figure 3-5: Trace of a Monitored Flow in an INSIGNIA System

3.5 TCP Performance

Most performance comparisons of MANET routing protocols have been conducted using UDP for the transport of constant bit rate traffic. However, TCP may be the dominant transport in mobile ad hoc networks. 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 that has been 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 performances of various TCP protocols, namely TCP-Reno [33], TCP-SACK [34], and TCP-Vegas [33] over the best-effort and INSIGNIA systems. We also evaluate the Explicit Link Failure Notification (ELFN) [35], which is specifically designed to enhance TCP in mobile ad hoc networks.

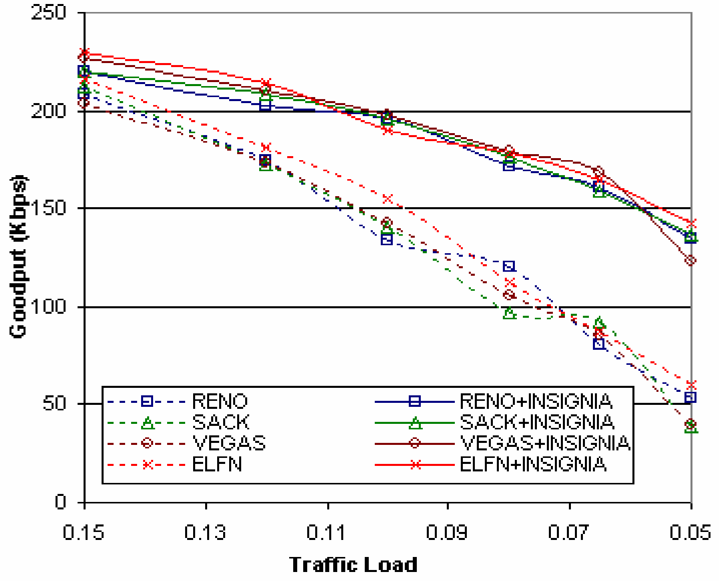

Figure 3-6(a): Impact of Network Traffic Load on TCP Goodput for Best-Effort and INSIGNIA Systems

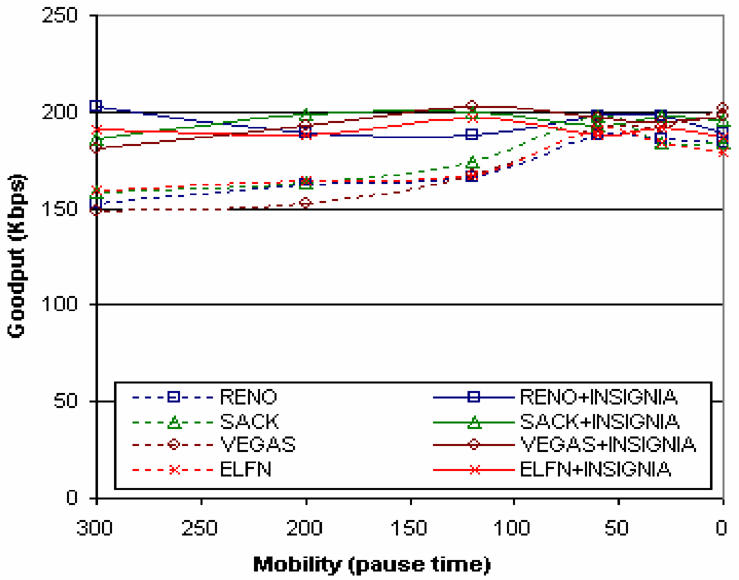

Figure 3-6(b): Impact of Mobility on TCP Goodput for Best-Effort and INSIGNIA Systems

3.5.1 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 sec pause time). The impact of increasing traffic load on TCP-Reno, TCP-SACK, TCP-Vegas and TCP-ELFN shows similar trends, as shown in Figure 3-6(a). The INSIGNIA system provides marginal improvement in goodput over the best-effort system when the network load is 628 Kbps (represented by packet generation interval of 0.15 sec). However, as the network load increases the performance improvement increases, as shown in the Figure 3-6(a). The goodput performance of the monitored flows decreases below 70 Kbps when the traffic load increases to 1.39 Mbps (maximum load) in the best-effort system. In contrast, the goodput of the six monitored flows in the INSIGNIA system remains above 125 Kbps under maximum load. This performance improvement represents a 150% increase in goodput for all versions of TCPs operating at maximum load. All TCP variants operate with some differentiation, as shown in the figure.

3.5.2 Impact of Mobility

The impact of mobility on TCP flows in terms of goodput is shown in Figure 3-6(b). To observe the impact of mobility on TCP goodput, we fix the traffic load at 800 Kbps and gradually increase the mobility of nodes. A traffic load of 800 Kbps is sufficient to produce congested conditions for the shared 2 Mbps 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 flow are within each other's transmission range the maximum available resources perceived by each mobile host is well below 1 Mbps. 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 Section 3.4. 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 re-routing 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 re-routing 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.

3.6 Conclusion

In this chapter, we have presented an overview of the INSIGNIA signaling system and evaluated the performance of AODV, DSR and TORA to operate in best effort and INSIGNIA systems. Furthermore, we have discussed 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.

Through extensive simulations and testbed implementation, we have shown that INSIGNIA provides substantial performance improvements to TCP and UDP sessions. The INSIGNIA ns-2 code used for the study reported in this chapter and actual testbed code for Linux platform are publicly available from our project website [71] (i.e., http://www.comet.columbia.edu/insignia).

While investigating the adaptive QOS issues presented in this chapter, we analyzed the performance degradations characteristics in our implementations (i.e., simulation and testbed) and observed that the encountered problems are mainly due to peculiar congestion conditions in MANETs. These congestion conditions were often transient but entailed significant packet loss, delay-spikes, and biased resource consumption. In the next chapter, we address this challenge and propose the first generic mechanism called HMP (Hotspot Mitigation Protocol) to mitigate the problem.

Chapter 4

A Hotspot Mitigation Protocol for Ad hoc Networks

4.1 Introduction

Hotspots are often created in regions of mobile ad hoc networks (MANETs) where flows converge and intersect with each other. We define hotspots as nodes that experience flash congestion conditions or excessive contention over longer time-scales (e.g., order of seconds). Under such conditions nodes typically consume more resources (e.g., energy) and attempt to receive, process, and forward packets but the performance of the packet forwarding and signaling functions is considerably diminished and limited during hotspot periods. This is the result of excessive contention of the shared media wireless access, and due to flash loading at hotspot nodes, and importantly, at neighboring nodes that are in the region of hotspots. Hotspots are often transient in nature because the mobility of nodes in the network continuously creates, removes, and to some degree, migrates hotspots because node mobility changes the network topology and causes flows to be rerouted. Hotspots are characterized by excessive contention, congestion, and resource exhaustion in these networks. In other words, hotspots appear when excessive contention exists, prompting congestion when insufficient resources are available to handle the increased traffic load.

Hotspots are intrinsic to many on-demand MANET routing protocols because most on-demand routing protocols [30] [31] utilize shortest path (or hop count) as their primary route creation metric. Most on-demand routing protocols allow an intermediate node to reply to a route query using cached route information, causing traffic to concentrate at certain nodes. We observe from our analysis of hotspots presented in this chapter that although many on-demand routing protocols prove to be effective in routing packets in these networks they also have a propensity to create hotspots. Other researchers have also made such observations [10][52][73]. We also observe that hotspot nodes consume a disproportionate amount of resources (e.g., energy).

In this chapter, we present a simple, effective, and scalable *Hotspot Mitigation Protocol (HMP),* which seamlessly operates with existing on-demand (e.g., AODV [30] and DSR [31]) and proactive (e.g., DSDV [36] and OLSR [37]) ad hoc routing protocols. HMP balances resource consumption among neighboring nodes and improves end-to-end throughput, delay, and packet loss. Our results indicate that HMP can also improve network connectivity preventing premature network partitions. Ideally, establishing routes through non-congested areas of the network and rerouting active flows away from congested areas to non-congested areas would be the best approach to hotspot mitigation. However, this requires extensive collaboration between nodes to establish load-aware routes and sophisticated algorithms to update time-varying loading conditions. Such an approach is unscalable and not practical in mobile ad hoc networks.

HMP represents a fully distributed and scalable protocol where nodes independently monitor local conditions and take local actions:

- *to declare* a node to be a hotspot if a combination of MAC contention/delays, packet loss, buffer occupancy, and remaining energy reserves exceed certain predefined system thresholds;
- *to suppress* new route requests at hotspots to ensure that routed traffic does not compound congestion problems; and
- *to throttle* traffic locally at hotspots to force TCP flows to slow down.

HMP also seeks to decrease the energy consumption of nodes in ad hoc networks via use of these mechanisms.

This chapter is structured as follows. In Section 4.2, we first analyze the behavior of hotspots using existing on-demand MANET routing protocols. Observations from this evaluation indicate that hotspots are evident even under relatively lightly loaded conditions in ad hoc networks, motivating the need for hotspot mitigation protocols. Related work is discussed in Section 4.3, followed by the design of the protocol in Section 4.4. We present a detailed analysis of HMP in Section 4.5 using both on-demand and proactive routing protocols and discuss results from the implementation of HMP in a wireless testbed in Section 4.6. Finally, in Section 4.7, we present some concluding remarks.

4.2 Hotspots

4.2.1 Existence of Hotspots

Hotspots are generally created when traffic converges to a node or small cluster of nodes. Flows traversing multiple wireless hops from various locations intersect with each other and create transient hotspot conditions. We observe that hotspot nodes and nodes in the vicinity of hotspots (i.e., in hotspot regions) are prone to consume more resources than others. Left unchecked such unbalanced resource consumption is detrimental to mobile ad hoc networks because overtaxed nodes would prematurely exhaust their energy reserves before other nodes. As a consequence the network connectivity can be unnecessarily impacted. In addition, we observe that hotspot nodes are often responsible for generating a large amount of routing overhead. In general, as the traffic load increases more hotspots appear and conditions in hotspot regions become aggravated.

In what follows, we make a number of observations about hotspots using ns-2 [40] and AODV [30]. Note that our observations are common to other on-demand protocols such as DSR [31]. The simulation consists of 100 mobile nodes in a 1200m by 1200m network under moderate mobility conditions (i.e., pause time of 80 seconds using the random waypoint mobility model with maximum speed of 10 m/sec). Thirty CBR/UDP and 10 TCP flows are used to produce an offered load of approximately 480 Kbps. We detect hotspots through a combination of MAC-delay measurements of unicast packets, packet loss, buffer occupancy, and by optionally considering the remaining energy reserves at a node. While the thresholds for these hotspot metrics are configurable, we considers a node to be a hotspot in our current implementation (which is based on IEEE 802.11), when the node consecutively measures *i)* MAC delays that exceed a predefined value, *ii)* packet loss during the RTS-CTS-DATA-ACK cycle, and *iii)* buffer overflow. We discuss these metrics and their configuration in Section 4.5 on hotspot detection.

Hotspots are often transient because of the mobility of nodes changes the topology and continuously varies the traffic load distributed across the network. We observe in our simulations that nodes are rarely in a permanent hotspot state. As a rule of thumb in our experimentation once a node is declared a hotspot it is marked as a hotspot for the next 5 seconds. Thus, under simulation, nodes could be declared a hotspot a number of times (e.g., 20 times) during the lifetime of the simulation run. Using this time-scale, we observe an average of 816 congestion hotspot incidents during a 300 second simulation described above where the offered load is 480 Kbps. Note, that 816 hotspots instances correspond to 4080 seconds of hotspot conditions in the network, or, an average of 40.8 seconds of hotspot conditions per node. Results are from 5 simulation runs.

4.2.2 Traffic Load

Figure 4-1 shows the packet delivery ratio (PDR), number of hotspots, and offered load for the simulation discussed above. The packet delivery ratio is defined as the total number of packets received out of the total number of packets sent. The offered load is varied from 50 Kbps to 963 Kbps under moderate mobility conditions involving 4831 link changes and 39830 route changes. The y-axis represents the packet delivery ratio and x-axis the offered load. In Figure 4-2, we also show the corresponding number of hotspot instances.

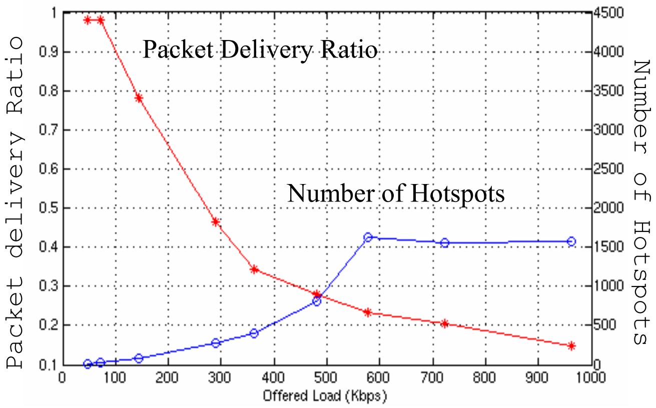

Figure 4-1: Packet Delivery Ratio and Number of Hotspots

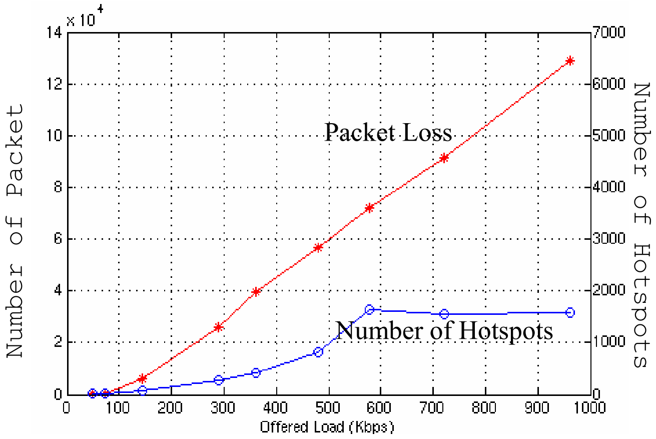

Figure 4-2: Packet Loss and the Number of Hotspots

As expected, the number of hotspots increases with offered load, while the packet delivery ratio decreases with increasing load. When the offered load is light, only few hotspots are detected where the network encounters few problems in routing packets. For example, when the traffic load is 72.2 Kbps, approximately 98% of packets are delivered correctly, and only 22 hotspot instances are detected during the simulation. This means that mobile nodes in the network encounter 110 seconds of congested conditions that in turn represents an average of 1.1 seconds/node of congestion. Note that link/route errors can occasionally be interpreted as congestion conditions because packet loss due to congestion is indiscernible from packet loss due to route failure.

When the offered load increases to 963 Kbps then only 15 % of the data packets are correctly delivered with 1566 hotspots instances observed. The difference is more than 70-fold when compared to an offered load of 72.2 Kbps. One interesting observation shown in Figure 4-2 is that number of hotspots levels-off when the offered load exceeds 580 Kbps. We identified that the reason for this anomaly is mainly due to the failure of congestion detection. All types of packets continuously fail to complete the collision avoidance cycle of IEEE 802.11 [50], and as a consequence, they are considered to be route errors while our hotspot detection mechanism, which relies on the measurement of the RTS-CTS-DATA-ACK cycle, fails to capture the congestion implications. The corresponding packet loss count observed during the simulation clearly supports this.

4.2.3 Overhead

Figure 4-3 illustrates the total number of packets transported when the offered load is 290 Kbps. The x-axis represents the node IDs and y-axis the number of packets handled by each node. Figure 4-3 also shows the number of data packets handled or forwarded by each node. One interesting observation is that most of the packets handled in the system are routing-related packets and only a small portion of the total transit traffic are data packets. For example, mobile node 2 handles 20103 packets in total during the simulation but only 1076 are data packets while 19027 are routing packets. Such observations are consistently observed in the network with the result that the ratio of signaling to data packets grows with the offered load.

The increase in the offered load aggravates congested conditions and as a consequence more packet loss is observed. Consecutive packet loss is often treated as route failures by ADOV triggering route recovery procedures that entail additional route requests, route errors, and route reply packet exchanges. It is observed that the routing overhead and number of hotspots increases with the offered load but begins to decrease beyond a certain load (e.g., 700 Kbps in this simulation set) due to substantial packet loss, as discussed earlier (i.e., route request packets continuously fail to be forwarded and rarely reach destination nodes, route replies are rarely generated, with the result that routes are seldom successfully established).

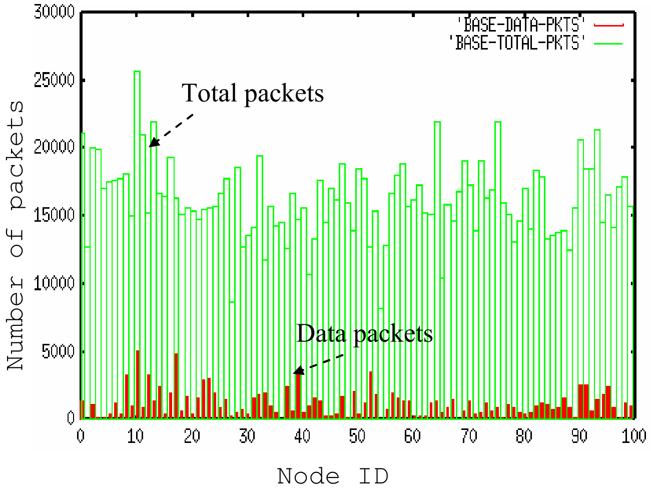

Figure 4-3: Packets Handled by Nodes

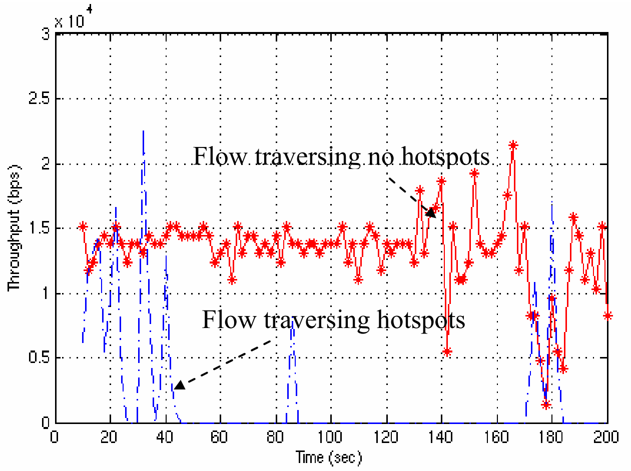

Figure 4-4: Throughput Traces of Two Monitored Flows

4.2.4 Hotspot Regions

Figure 4-4 shows the throughput traces of two similar flows under the simulation configuration discussed previously. We selected a flow traversing multiple hotspots and a flow encountering no hotspots (from our simulation results) and compare their throughput performance. The trace intuitively demonstrates how hotspots impact flow performance. Among 100 mobile nodes, 11 nodes are identified as severe hotspots where they experience congestion for more than 110 seconds out of the 200 seconds monitoring period. We identified 59 nodes as immediate neighbors of the 11 severe hotspots. We observed the packet loss of these 70 nodes (i.e., 11 severe hotspots and their 59 neighbors) that resided in hotspot regions, and compared their performance to other nodes in the rest of the network. We observed that nodes residing close to hotspot nodes also experience degradation in performance. For example, when the offered load is 290 Kbps, hotspot regions are responsible for 94.9 % of total packet loss while the rest of the network contributed only 5.1 % to the total packet loss. Moreover, nodes in hotspot regions have an average congestion time of 94 seconds while the rest of the network nodes only experience 36 seconds of congestion time. Based on these observations we argue that there is a need to study, design, and evaluate mechanisms that can seamlessly interwork with existing routing protocols to mitigate the impact of hotspots in MANETs.

4.3 Related Work

MANET routing protocols can be simply classified into best effort routing protocols that have no built in mechanisms to provide better than best effort service [30] [31] [32], QOS-based routing protocols [75] [76] [77], and multipath routing protocols [77] [82]. While HMP is not a routing protocol it is designed to interwork with the existing best effort routing protocols (e.g., on-demand and proactive protocols) to provide hotspot mitigation support.

Currently, none of the existing on-demand best effort routing protocols take hotspots into account in their routing decisions. As shown in the last section this allows hotspots to quickly emerge and build up in the network under normal operating conditions. There is a clear need to propose new mechanisms that can interwork with, or be directly incorporated into, these best effort routing protocols, hence enhancing the network's performance. HMP is designed as a separate mechanism and is therefore capable of being used in combination with any of the existing best effort routing schemes.

HMP incorporates measures of congestion and contention as well as resource shortages (e.g., energy) into its definition of hotspots. We believe that this is a more realistic definition for wireless mobile networks than one that only considers the buffer occupancy statistics at intermediate nodes. Using buffer occupancy as an indication of congestion has been widely used by a number of Internet congestion control/ hotspot management schemes. HMP manages these hotspots locally (i.e., at the point of interest) in a fully distributed fashion, as opposed to traditional end-to-end approach for managing congestion.

The simple goal of HMP is to disperse new flows away from being routed through hotspots and congestion-prone areas (i.e., hotspot regions), avoiding the further build up of traffic load at hotspots or in hotspot regions. HMP distinguishes itself from the various QOS routing approaches, which in practice are complex to implement, in that

HMP does not attempt to provide QOS support nor QOS routes. However, the deployment of QOS routing and multipath routing algorithms would also minimize the likelihood of hotspots, but not eradicate them. QOS routing algorithms require accurate link state (e.g., available bandwidth, packet loss rate, estimated delay, etc.) but due to the time-varying capacity of wireless links, limited resources and mobility, maintaining accurate state information is very difficult if not impossible in mobile ad hoc networks. Finding a feasible route with just two independent path constraints is an NP-complete problem [83]. Moreover, finding a QOS satisfying path is merely the first part of the problem because it is more challenging to maintain QOS routes when the network topology changes [99]. Because QOS routing relies on this distributed but global review of resources in the network the likelihood of stale state and traffic fluctuations beyond the anticipated load also calls for localized reactive mechanisms such as HMP to help alleviate transient hotspots. We therefore consider that HMP would also be useful in QOS routed networks.

Alternate path routing and multipath routing protocols can outperform single path routing protocols. A common feature of these protocols is that they utilize backup or alternate routes when primary routes fail. Some multipath routing protocols are designed to distribute traffic among multiple paths and reassemble the traffic at the destination nodes. However, reassembling traffic at the destination node in this manner can be problematic because it leads to out-of-sequence delivery and extra re-sequencing delays. Moreover, maintaining additional path information requires additional routing and computational overhead. Alternate paths should be comprised of disjoint-paths [82] in order to be effective. Such alternate paths often do not exist, particularly in single channel ad hoc networks (e.g., based on IEEE 802.11).

In summary, HMP is designed as a localized node mechanism that takes local actions to prevent the build up of hotspots, which we believe will be very likely in MANETs under normal operating conditions. While HMP is targeted to interwork with the existing best effort routing protocols it could also provide efficient support for hotspot mitigation in MANET networks based on QOS routing and multipath routing. This is the subject of future work.

4.4 Hotspot Mitigation Protocol

4.4.1 Protocol Operations

The main goal of HMP is to redirect new “routes” away from hotspots. HMP disperses new flows away from being routed through hotspots and congestion-prone areas, avoiding the further build up of traffic load in hotspot regions. HMP effectively mitigates hotspot conditions and reduces congestion-related problems. Mitigating hotspot in this manner also helps to balance the resource consumption among neighboring nodes, and can extend the lifetime of overtaxed nodes.

The protocol utilizes MAC-delay measurements, packet loss, buffer occupancy information, neighbor status information and other resource monitoring mechanisms (i.e., energy) to detect hotspots. HMP does not limit the scope of monitoring and detection mechanisms, however. Operators are free to introduce additional mechanisms and algorithms according to their needs. In fact, we envision that a HMP network would embody diverse mechanisms operating concurrently. HMP utilizes measured information to respond to conditions by executing the most appropriate algorithms to alleviate the condition at hand. The measured conditions are expressed by a multimetric parameter called STATUS, which consists of two components: *symptom* and *severity*. Symptom describes the dominant condition a node is experiencing while severity expresses the degree of the symptom. For example, a node may declare its status as $Y_{CONGESTION}$ while another node may declare its status as R_{ENERGY}. This status is analogous to traffic lights, where green (denoted by G) indicates a good condition, yellow (Y) represents a marginal condition, and red (R) represents a critical condition. Therefore, $Y_{CONGESTION}$ indicates marginal congestion and R_{ENERGY} indicates critically low energy reserves. Users/operators are free to introduce more granularity if needed. HMP piggybacks this status information in the IP option field and neighboring nodes operating in promiscuous mode learn the status of transmitters by eavesdropping their packets. The eavesdropped information is used to create and update a *Neighborhood Status Table (NST).* This status information is cached and locally maintained and updated at each node.

A node's NST caches a list of immediate neighbors and their status. It is primarily used to manipulate new-route-creation decisions at nodes. In other words, a node refers to its NST to ensure that it is not aggravating the conditions of neighboring nodes by creating new routes through them. We assume a finite number of neighboring nodes surrounding any node, which in effect defines the size of the NST at a node.

The naïve suppression of new route creation may prevent the use of the only possible path between two hosts and may yield poor connectivity in the network, or even cause network partitions. To avoid this, a new-route-suppression mechanism is used, if and only if, there exists a sufficient number of non-hotspot neighbors within its transmission range. HMP also makes sure that preceding nodes en-route also have enough non-hotspot neighbors. The notion of 'enough neighbors' is defined by the $NUM_{ENOUGH\text{-}NEIGHBOR}$ parameter. This is currently set to 6 in the testbed implementation discussed in Section 4.6. The value of this parameter has a direct impact on the network connectivity, as discussed in Section 4.5. If $NUM_{ENOUGH\text{-}NEIGHBOR}$ is too small, (e.g., 2), then HMP manifests low connectivity among mobile nodes and often fails to provide useful routes. HMP also ensures that it is not inadvertently denying the only possible path between two end hosts by utilizing an indicator called the *path_indicator*, which is carried in the IP option field of Route Request (RREQ) messages. A node that has only a few neighbors sets this indicator (path_indicator = 1) and upstream nodes that receive the RREQ (with IP option that includes path_indicator) check this indicator and avoid suppressing new routes if it is set. This is illustrated in Figure 4-5 where hotspot node M_4 forwards RREQ toward node M_5 because the source node M_3 has set its path indicator whereas hotspot M_2 suppresses the RREQ message from M_1 because its path_indicator is not set in the IP option field of the RREQ.

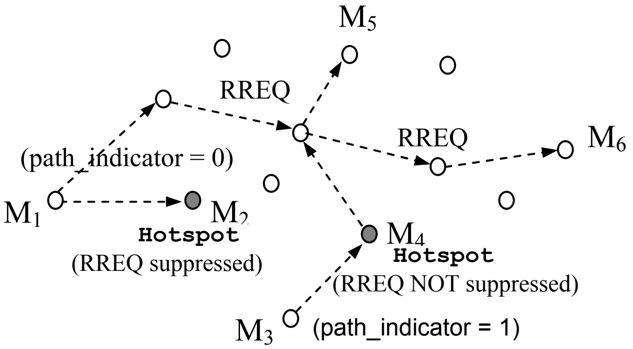

Figure 4-5: Hotspot Mitigation Protocol Illustration

4.4.2 Congestion Levels

The main objective of congestion avoidance algorithms is preventing the further build up of traffic at hotspots. HMP distinguishes two levels of congestion (i.e., levels 1 and 2) and adopts two corresponding algorithms to support this view. The first algorithm is activated when HMP determines the current status of a node is in a moderately congested condition (i.e., level 1), denoted by $Y_{\text{CONGESTION-1}}$. This algorithm simply suppresses the creation of additional routes at hotspots by discarding new route request packets. As mentioned previously, HMP ensures not to deny the 'only route' between two hosts.

The second algorithm is more aggressive and executes when nodes encounter substantial congestion (i.e., level 2), denoted by $Y_{\text{CONGESTION-2}}$. This algorithm is executed when a node experiences severe hotspot conditions without any non-hotspot neighbors. This algorithm not only suppresses new route creation but also throttles best effort TCP flows traversing the node in an attempt to reduce the load using rate control mechanisms discussed in [78]. TCP flows are bandwidth hungry and unless controlled can easily occupy all remaining wireless medium bandwidth. Throttling TCP rates locally in this manner does not necessarily hurt TCP sessions but can effectively relieve congestion bottlenecks. Users and operators are free to introduce other schemes to relieve congestion conditions, e.g., one simple policy is dropping TCP packets at bottleneck nodes.

HMP attacks the congestion at the point of congestion (POC) as opposed to a traditional end-to-end approach. Although congestion is an end-to-end issue where it is detected and controlled (e.g., as in the case of TCP), traditional remedies for end-to-end congestion control are not effective in mobile ad hoc networks. In fact, such traditional control mechanisms may limit the utilization of the wireless medium that is constrained by hotspots. We argue that we can avoid such shortcomings if we tackle the problem at the point of congestion rather than responding on end-to-end basis.

4.4.3 Hotspot Detection

A number of system parameters are used by HMP to identify hotspots. These per-node parameters, which are associated with MAC-delays, packet loss, and buffer occupancy, can be configured to make HMP's hotspot detection mechanism more or less aggressive in its declaration of hotspots.

The $M_{DELAY\text{-}THRESH}$ parameter is used by the protocol to detect MAC-delay violations. If the measured MAC-delay exceeds the $M_{DELAY\text{-}THRESH}$ then HMP considers this a MAC-delay violation. If the number of these violations exceeds a predefined value called N_{THRESH} then the protocol takes a number of actions discussed below. We define the MAC-delay as the measured time for the successful transmission of a data packet at the MAC layer. This includes the time taken for the RTS-CTS-DATA-ACK message exchange over the air. Because IEEE 802.11 defines up to 7 possible retransmissions of a data packet the measured MAC-delay could represent up to a maximum of 7 RTS-CTS-DATA-ACK cycles in the case were packet loss occurs. Each node continuously monitors the on-going MAC-delay and compares it to the $M_{DELAY-THRESH}$ value, which is computed as the average of the minimum and maximum MAC-delays. The N_{THRESH} parameter is used to control the sensitivity of the protocol to the measured MAC-delays. N_{THRESH} defines how many consecutive MAC-delay violations can be tolerated before a node is declared a hotspot. This parameter essentially determines how aggressive HMP is in declaring hotspots. In other words, a node is identified as a hotspot when the $M_{DELAY-THRESH}$ parameter is consecutively violated more than N_{THRESH} times. Hotspot detection also needs to consider the case of packet loss too. In the case where there is an intermittent packet loss between two consecutive MAC-delay violations, HMP takes account of this condition during hotspot detection; that is, a node is also considered a hotspot when the $M_{DELAY-THRESH}$ parameter is violated $N_{THRESH} – (\varphi)$ times, where φ is defined as the number of intermittent packet losses during a hotspot detection interval. A hotspot detection interval starts when the first MAC delay violation is observed and lasts until the node either declares itself a hotspot based on the criteria described above or a data packet is successfully delivered without a MAC-delay violation. The MAC-delay violation count maintained by HMP during the hotspot detection interval is reset at the beginning of a new interval. Note that many MANET routing protocols consider that three consecutive packet losses represents link or route failure. Any link failure or route error also resets all associated counters/parameters used by HMP's hotspot detection.

HMP also monitors buffer occupancy to identify hotspots. If a node detects that the buffer occupancy exceeds a predefined threshold called B_{THRESH} then it will check for MAC-delay violations. The B_{THRESH} parameter is set to a buffer level that is less than the buffer overflow mark. If B_{THRESH} is exceeded and there is at least one MAC-delay violation then the node declares itself a hotspot. We adopt this hybrid approach to hotspot detection because buffer occupancy information alone is insufficient to declare a hotspot unless the buffer overflow mark is exceeded. As a result we combine buffer occupancy with MAC-delay violations to make the approach more accurate.

In summary, hotspots are declared by HMP if MAC-delays and packet loss violate a predefined threshold, or buffer occupancy exceeds a given level and at least one MAC-delay violation is observed, or when the buffer occupancy exceeds the buffer overflow mark.

4.4.4 Energy Conservation

Mobile ad hoc networks are essentially energy-limited networks and are likely to be comprised of heterogeneous nodes with diverse energy constraints. Some mobile devices will have large energy reserves in comparison to others. There exist various energy-aware power-conserving protocols for mobile ad hoc networks [84]. The common objective of these protocols lie in conserving energy as much as possible to prolong the lifetime of the network or extend the lifetime of individual nodes.

Although energy conservation is not a primary concern of HMP, the protocol provides a simple mechanism to conserve energy through its status declaration mechanism. A node with limited energy reserves can declare itself a hotspot by setting its status to Y_{ENERGY} or R_{ENERGY} when its energy reserves are marginally or critically low, respectively. The triggering thresholds are $P_{YELLOW\text{-}THRESH}$ and $P_{RED\text{-}THRESH}$. In our current implementation, $P_{YELLOW\text{-}THRESH}$ is set to 50% of node's initial (or maximum) energy reserves and $P_{RED\text{-}THRESH}$ is fixed at 1.00 joule. The latter value represents the amount of energy needed for a node to sustain a CBR flow for approximately 300 packets in most of our simulation sets. However, we note that operators and users are free to set these values according to their own needs, based on the characteristics of the targeted network. A node with energy concerns is acknowledged by neighboring nodes and new route creation through such a node is avoided if possible. On the other hand, a node with critical energy (i.e., R_{ENERGY} status) immediately relinquishes its role as a router and functions strictly as an end host in order to conserve energy (maximize its lifetime) unless it is identified as the only intermediate node between two communicating end hosts.

4.5 Performance Evaluation

In what follows, we evaluate HMP using simulation and discuss the performance improvements that the protocol offers. Simulation metrics such as packet delivery ratio, packet loss, throughput, end-to-end delay, per-hop delay, and energy consumption are used in the evaluation of the protocol. We also discuss the impact of various parameters on the performance of HMP. In the initial part of the evaluation we use the AODV [30] routing protocol with HMP, and in the latter part, DSR [31] with HMP. In addition to discussing on-demand routing protocols we also discuss the performance of HMP with proactive routing schemes including DSDV [36] and OLSR [37].

We implemented HMP using the ns-2 simulator and its wireless extension. The HMP implementation includes monitoring modules, measurement mechanisms, an NST module, and the HMP algorithms discussed in Section 4.4. The simulated network size is 1200 meters by 1200 meters where 100 mobile nodes create 10 TCP and 30 CBR/UDP flows that arbitrarily last for 60 to 280 seconds. Moderate mobility is assumed with a pause time of 80 second using the random way point mobility model [30] [31] unless specified otherwise. All data packets are a fixed size of 128 bytes, each simulation run lasts for 300 seconds, and each data point represents an average of 5 simulation runs with the identical traffic model but different mobility scenarios. Each mobile node has a transmission range of 250 meters and shares a 2 Mbps radio channel with its neighboring nodes. The simulations also include a two-ray ground reflection model, finite energy module, and the IEEE 802.11 MAC protocol. Throughout the evaluation section we use the terms ‘HMP system’ and ‘baseline system’ to refer to wireless ad-hoc networks with and without the HMP mechanisms, respectively.

4.5.1 Hotspot Detection Analysis

Accurate and timely hotspot detection is one of most crucial aspect of HMP. To determine hotspots, the protocol relies on MAC-delay measurements, packet loss detection in RTS-CTS-DATA-ACK exchanges, buffer occupancy, and, if selected, the residual node energy. Among the measurements we have observed that the MAC-delay measurement is the most useful indicator since a hotspot always manifest in increased delays in the RTS-CTS-DATA-ACK cycle. As stated earlier, relying solely on the buffer occupancy is rather inaccurate. We often witnessed that hotspot conditions are created without any buffer occupancy. Such events are due to excessive contention among neighboring nodes. We discuss such observation below. Therefore, in order to minimize the margin of error in hotspot detection, we utilize both buffer information and MAC-delay measurements together with some other additional system parameters discussed later.

Figure 4-6 shows a typical trace of the MAC-delay measurement of a node. The x-axis represents the simulation time and y-axis represents the MAC-delay measurements of a randomly selected mobile node. As shown in the figure, MAC-delay measurements continuously fluctuate throughout the simulation. Spikes in the delay trace typically represent congested conditions, while zero delay measurements are observed when the node is not participating in the RTS-CTS-DATA-ACK activity. Recall that detection of a hotspot is dependent on two key parameters: (i) MAC-delay threshold (i.e., denoted by $M_{DELAY\text{-}THRESH}$), which determines when a packet is considered a delayed packet; and (ii) N_{THRESH}, which determines when a node is considered a hotspot. Specifically, a node is considered a hotspot when the measured MAC-delay exceeds a predetermined threshold (i.e., $M_{DELAY\text{-}THRESH}$) for more than N_{THRESH} consecutive times. These two parameters have an impact on how many hotspots are detected by HMP. When $M_{DELAY\text{-}THRESH}$ and N_{THRESH} are configured as large values, HMP is too conservative and only detects a small number of hotspots rendering the protocol to be less effective against moderate congestion. In contrast, when $M_{DELAY\text{-}THRESH}$ and N_{THRESH} are configured with small values, HMP is aggressive and detects too many hotspots too hastily. Therefore, the appropriate choice of these parameters is important for HMP to function properly. The use of the N_{THRESH} parameter also prevents HMP from premature detection of a hotspot when experiencing a momentary increase (i.e., a spike) in the MAC-delay measurement. It was observed that the MAC-delay measurements intermittently “spike” without any noticeable congestion conditions (e.g., during rerouting). To avoid reacting to such transient behavior and to increase the accuracy of hotspot detection, HMP marks a node as a hotspot, if and only if, the MAC-delay measurements are violated (i.e., exceeds $M_{DELAY\text{-}THRESH}$) more than N_{THRESH} consecutive times. Currently, $M_{DELAY\text{-}THRESH}$ is set to 20 msec and N_{THRESH} is set to 4.

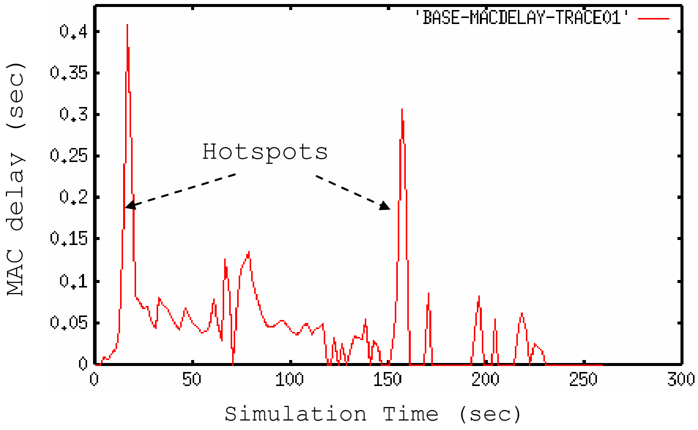

Figure 4-6: Trace of MAC-delay Measurements

We evaluate a number of different configurations of the protocol based on these parameters and studied the sensitivity of these parameters to HMP's ability to efficiently and accurately detect and mitigate hotspots in MANETs. We analyze the impact of buffer occupancy thresholds on HMP's hotspot detection using various levels of buffer occupancy thresholds. Note that other system parameters such as MAC-delay, packet loss, are omitted from the hotspot decision algorithm in order to solely monitor the impact of different buffer thresholds on HMP's performance. Results from three different buffer threshold settings are shown in Figure 4-7. We observe HMP performance when the buffer threshold is set to 10 %, 30%, and 60% of the total buffer capacity, respectively. Figure 4-7 also shows results of the baseline system and the original version of HMP that detects hotspots using a full set of hotspot detection parameters and not solely buffer occupancy.

As observed in Figure 4-7, sole use of buffer occupancy information is not an accurate measure to detect hotspots. A comparison of the packet delivery ratio shows that regardless of the choice of the B_{THRESH} (viz. 10%, 30%, 60%) value the original version of HMP (which uses MAC-delay and packet loss for hotspot detection) out performs the various versions of HMP configured to only use buffer occupancy as an indicator of congestion. However, that is not to say that buffer occupancy is not useful in combating congestion. Rather, buffer occupancy may be a better indicator if used in combination with information on the build up of delays or packet loss. Another interesting observation is all the versions of HMP that solely used buffer occupancy as an indicator of hotspots did better than the baseline system, which does not use any detection mechanism. This provides an interesting insight on hotspot mitigation because the outcome indicates that any ‘reasonable form’ of hotspot mitigation provides performance improvement over the baseline system.

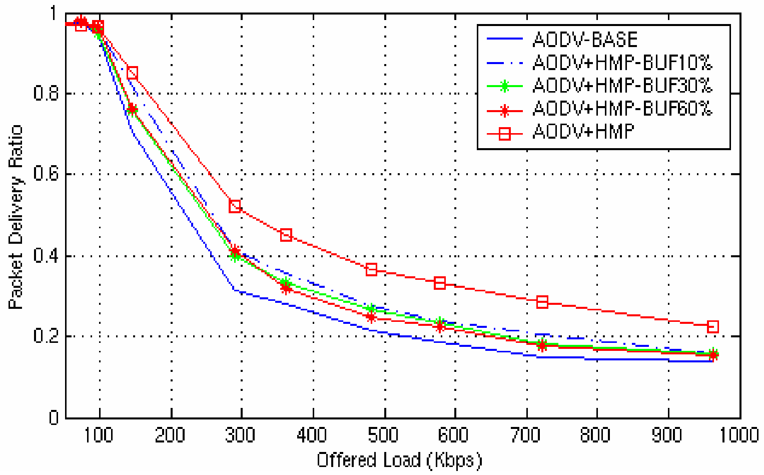

Figure 4-7: Impact of Buffer Occupancy Thresholds on Hotspot Detection

4.5.2 Throughput Analysis

We first observe how the HMP system performs in comparison to the baseline system in terms of the packet delivery ratio (PDR). Figure 4-8 shows a comparison of the packet delivery ratio against increasing load for two different HMP system configurations (discussed below) and the baseline system. The two HMP systems are simply called HMP-P and HMP-R where HMP-R is more aggressive than HMP-P in its route suppression mechanism. HMP-P stands for HMP-POC where HMP mechanisms are executed only at points of congestion (POC). On the other hand, HMP-R represents HMP-Regional signifying the regional execution of hotspot mitigation algorithms. In other words, when a hotspot is detected HMP-P executes hotspot mitigation algorithms at the point of hotspots whereas HMP-R executes its mechanisms across a hotspot region. A node belongs to a hotspot region if it is a hotspot or it is an immediate neighbor of a hotspot. We note that both $NUM_{ENOUGH-NEIGHBOR}$ and path_indicator are always considered in all hotspot mitigation decisions.

Figure 4-8 shows that HMP-P and HMP-R have little impact when operating in lightly loaded networks, (e.g., below 100 Kbps). This is because the baseline system already achieves more than 90 % PDR and HMP has little room to make any improvements.

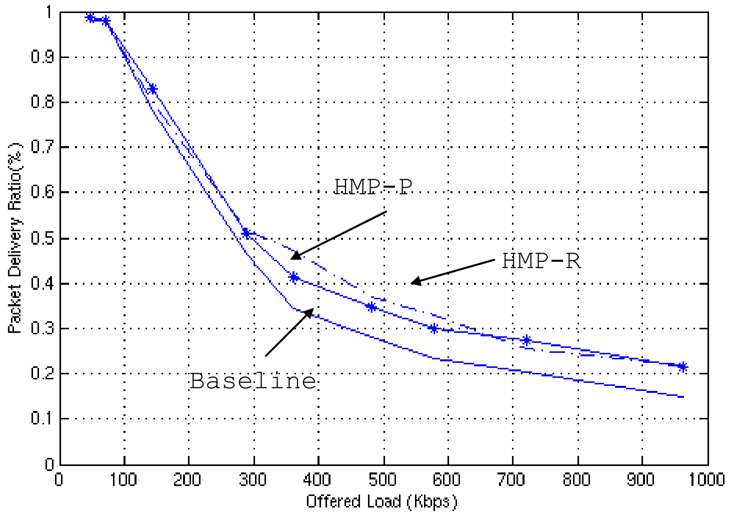

Figure 4-8: Comparison of PDR against network load

However, as the offered load increases, and congestion builds up, HMP begins to provide improvements, as shown in the figure. Both HMP-P and HMP-R provide substantial improvements in the PDR. Specifically, HMP-P and HMP-R provide up to a 43% and 46% increase in the packet delivery ratio when compared to the baseline system performance. From Figure 4-8, we also observe the behavior of HMP-R is more aggressive than that of HMP-P. When the offered load is moderately high, HMP-R often outperforms HMP-P and the baseline systems but becomes less effective when the offered load is light, (e.g., below 250 Kbps). The performance of HMP-R varies with different loads, as shown in Figure 4-8. We conclude that HMP-R is too aggressive for lightly loaded networks rendering it only useful in heavily loaded networks.

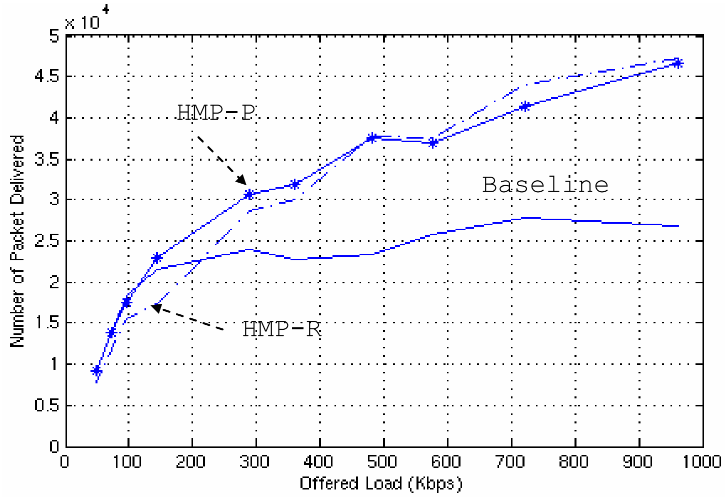

Figure 4-9: Number of Data Packets Delivered

Further analysis of the HMP-P, HMP-R and baseline systems can be seen by inspecting the number of delivered packets, as shown in Figure 4-9. An interesting observation is that number of packets delivered by the baseline system levels-off around 2.3 x 10^4 delivered packets but in the HMP-P and HMP-R systems the number of delivered packets continuously increases with increasing offered load. There are two major reasons for this improvement. First, HMP creates routes through non-congested nodes whenever possible allowing networks to utilize more distributed routes in the network even if these routes are not the shortest path. Creating routes at non-hotspot nodes allows traversing flows to encounter fewer problems, and as a consequence, more packets are delivered. Second, HMP generates less routing overhead when hotspots suppress new routes. Many hotspot nodes rebroadcast route request packets and these packets often flood large areas of the network or even the entire network. However, many of these rebroadcast route request packets are lost before reaching destination nodes. We observed that a considerable amount of route request packets are just wasted in the network without successful route creation in heavily loaded networks.

In the HMP systems, routing packets (i.e., route request) are pre-filtered at hotspot nodes/regions. This not only prevents new routes being created through hotspots but also helps reduces the number of wasted new route requested packets (that rely on broadcast/flooding), which are likely to be lost. This opens up room for more data packets, and as consequence, more packets are delivered in HMP systems in comparison to the baseline system. As congestion become more severe more nodes encounter packet loss and often interpret this packet loss as route errors, triggering route recovery routines. As a consequence, additional routing overhead is added to an already congested network. In HMP networks, congested nodes avoid participating in new route creation to mitigate congested conditions, and consequently less routing packets are observed in the network.

We observe that HMP-R outperforms HMP-P when the offered load is heavy. However, HMP-R is too aggressive for lightly loaded networks. We observe that the PDR of HMP-R is less than that of HMP-P and no better than that of the baseline system when the offered load is less than 150 Kbps. However, both HMP systems outperform the baseline system. In what follows, we refer to HMP-P when we discuss HMP unless otherwise stated.

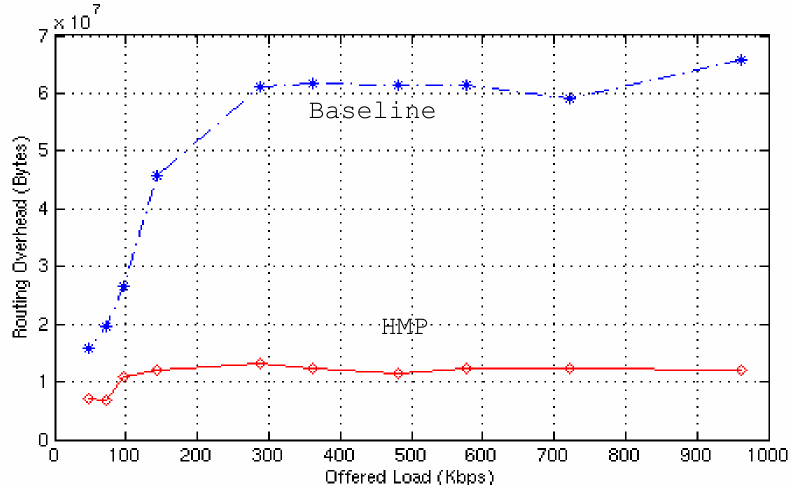

Figure 4-10: Comparison of Routing Overhead

4.5.3 Routing Overhead Analysis

Figure 4-10 shows the routing overhead of the HMP and baseline systems accumulated over 300 seconds of simulation time. The advantage of the HMP system over the baseline system in terms of the routing overhead is shown in the figure. It is observed that the HMP system provides up to a 75 % reduction in the routing overhead over the baseline system because of better route selection and routing packet suppression in hotspot regions. For example, when the offered load is 722 Kbps, the baseline system generates approximately 59 x 10^6 bytes of routing overhead while the HMP system only generates 13.3 x 10^6 bytes of routing overhead.

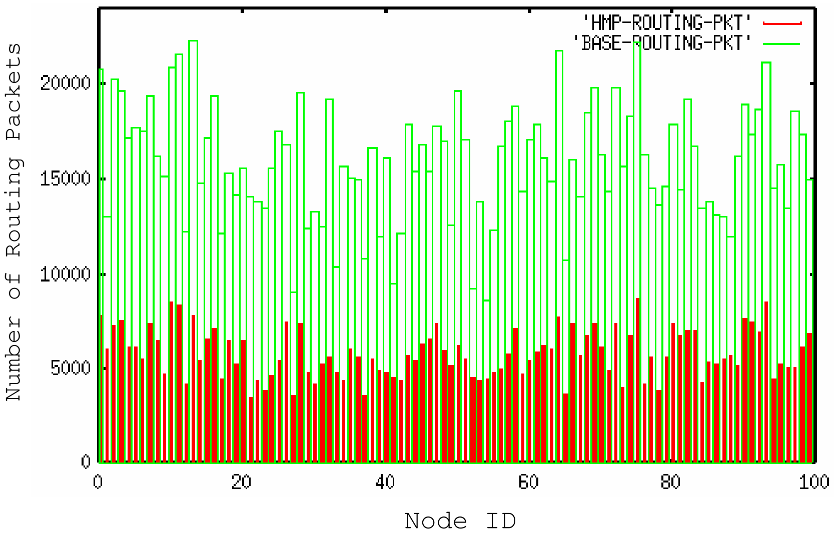

Figure 4-11: Routing Overhead Compared

Figure 4-11 compares the total number of routing packets transported by the HMP and baseline systems when the offered load is 577 Kbps. As expected there is a substantial difference between the two systems. It is observed that the HMP system always carries less routing load in comparison to the baseline system. This implies that HMP is not over-suppressing routes because if connectivity were limited, the number of route request packets would quickly increase and be reflected in the routing overhead. Therefore, it is safe to say that HMP provides sufficient connectivity under all the simulation scenarios. The HMP system outperforms the baseline system in terms of the PDR, number of packets delivered, and the routing overhead. These improvements are mainly due to effective hotspot mitigation through implicit route dispersal and suppression of new route request packets. HMP is prudent in route suppression decisions while ensuring sufficient connectivity when the configuration of the system (i.e., $M_{DELAY\text{-}THRESH}$, N_{THRESH}, $NUM_{ENOUGH\text{-}NEIGHBOR}$ and path_indicator) is enforced.

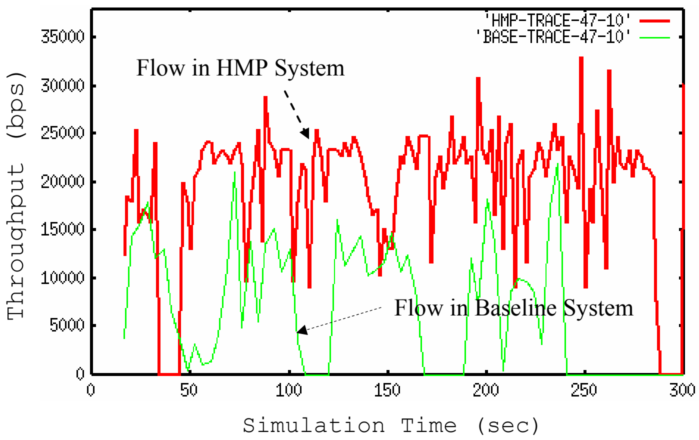

Figure 4-12: Throughput Trace Comparisons

Next, we compare throughput traces of a flow in the two systems. Figure 4-12 shows a monitored flow between node 47 and node 10. The monitored flow in the HMP system shows substantial improvements over the baseline system. More importantly, it is observed that the monitored flow traverses different routes in the two systems. Specifically, flow 47-10 traverses nodes 16, 18, 43, 51, 78 and 83 in the baseline system and traverses 16, 21, 38, 65, 78 and 83 in HMP system, during the monitored period of 50 seconds. Nodes 18 and 43 are identified as hotspots and consequently flow 47-10 avoids these two hotspots when using HMP. Such characteristics are consistently observed throughout the simulation, and as a consequence, the HMP system provides better throughput performance.

The previous evaluation of HMP considered AODV routing only. In what follows, we describe how HMP performs with DSR [31]. First, we observe the PDR traces of the baseline DSR system in comparison to the HMP+DSR system. Figure 4-13 shows the PDR trace for increasing offered load with moderate mobility for these systems. The figure also includes the PDR trace for the baseline AODV system and the AODV+HMP system (taken from Section 4.5) for comparison purposes.

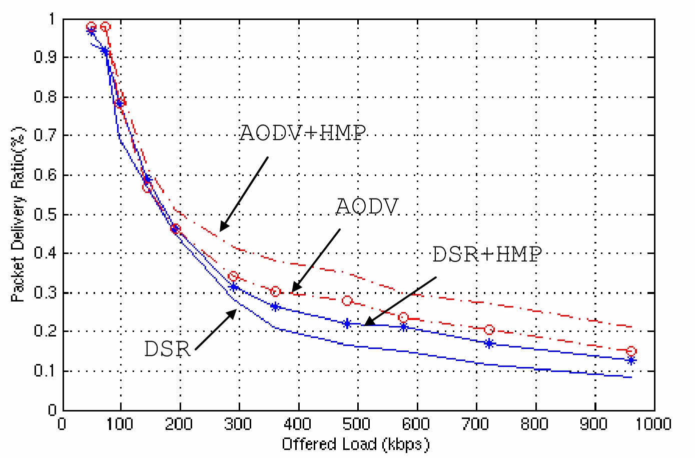

Figure 4-13: Impact of HMP on DSR and AODV

As expected, the DSR+HMP system provides improvements over the baseline DSR system. From Figure 4-13, it is observed that all the systems demonstrate similar performances under lightly loaded conditions but they begin to diverge as the offered load increases. One interesting observation is that DSR and AODV display different performance against offered load. They show similar results only under lightly loaded conditions. As congestion intensifies AODV begins to outperform DSR. This observation coincides with the results reported in [52]. We can observe from Figure 4-14 that HMP provides substantial reductions in routing overhead when operating with AODV but demonstrates different results with DSR. The main reason for this is related to the amount of routing load reduction. The difference in routing overhead reduction is directly reflected in the PDR traces (or when we compare the number of packets delivered). HMP provides improvements with AODV mainly through the reduction in the routing overhead, and route diversion away from hotspots. In contrast, in the DSR system the dominant reason for improvement is mainly due to route diversion from hotspots. We also observe that DSR's aggressive use of route-cache limits its performance too; that is, under harsh conditions (i.e., increased mobility, increased load), it can be observed that DSR maintains stale routes, generating a large amount of route-error messages. This observation is also reported in [52]. HMP successfully routes traffic through non-hotspot nodes but DSR's route-optimization scheme [31] utilizes cached routes, which often introduce new hotspots. Figure 4-14 reflects this observation.

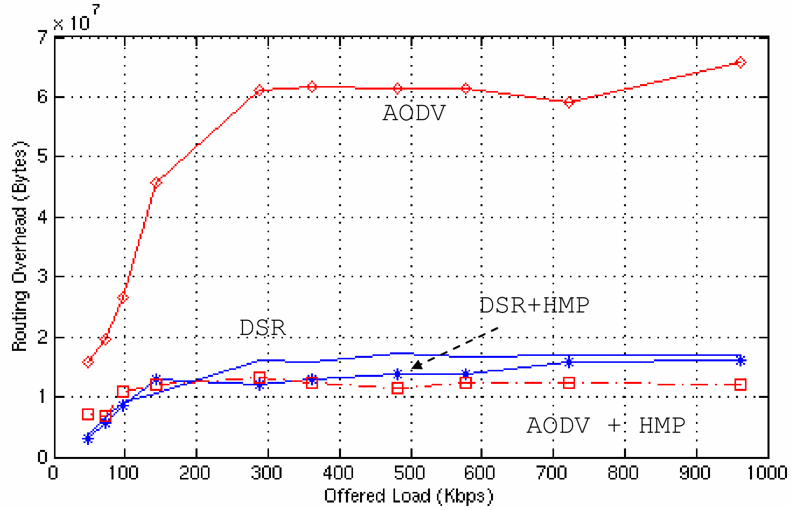

Figure 4-14: Routing Overhead Compared

4.5.4 Energy Analysis

We note that energy consumption is concentrated in hotspot regions and nodes. Figure 4-15 shows measurements of residual energy for nodes at the end of the simulation run. We assign a uniform energy of 25 joules to each node and conducted simulations for 100 seconds with AODV. The x-axis represents node IDs and y-axis represents the residual energy in joules. Bars represent the residual energy measurements of the baseline system and the superimposed impulses represent the corresponding measurements of the HMP system. As shown in the plot, the energy conservation provided by HMP for nodes in hotspot regions is significant. The baseline system exhibits 21 energy-depleted nodes (i.e., remaining energy is less than 0.01 joule such that it can no longer participate in communications) while there is not even one depleted node in the HMP system at *t* = 100 sec. Note that HMP improves packet delivery ratio, delay measurements, and reduces routing overheads, while providing energy conservation in hotspot regions.

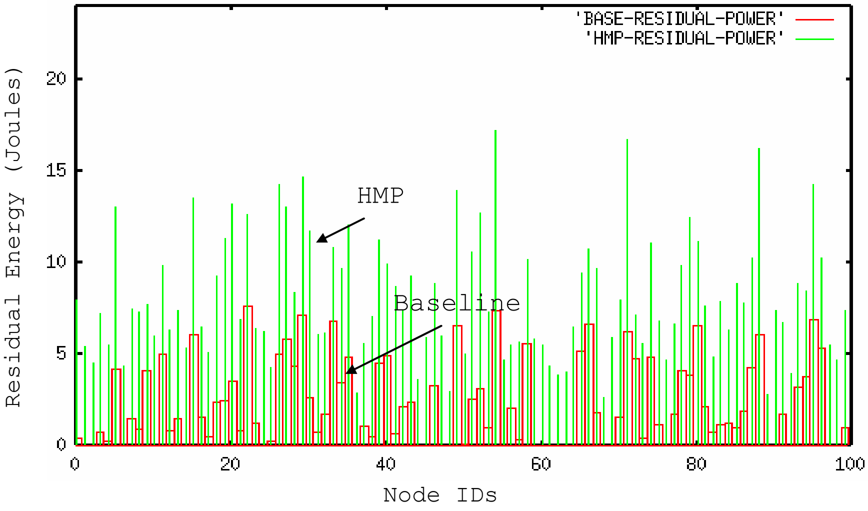

Figure 4-15: Residual Energy Compared

4.5.5 Sensitivity Analysis

In what follows, we describe four different HMP system configurations to study the responsiveness of the protocol to detect and mitigate hotspots. Four key parameters govern the HMP system control mechanisms; these are, $M_{DELAY-THRESH}$, N_{THRESH}, $NUM_{ENOUGH-NEIGHBOR}$ and path_indicator. For example, if the $M_{DELAY-THRESH}$ value is too small HMP may become too aggressive and declare too many hotspots. A small increase in the MAC-delay threshold measurement (or jitter) may falsely be recognized as congestion with many nodes being claimed as hotspots. In contrast, if the $M_{DELAY-THRESH}$ value is too large HMP may not identify any hotspots in the network and relegate itself to the baseline system. The second parameter N_{THRESH} is used to prevent

HMP from reacting to transient behavior. A momentary increase in the MAC-delay measurement and buffer occupancy are not necessarily a product of congestion or excessive contention. Delay may be observed for a very short period due to the rerouting of flows or a small burst of route query packets. Reacting to such transitory phenomenon is not beneficial because real hotspots cannot be distinguished from transient events. The third parameter is the path_indicator, which indicates that insufficient conditions exist for new route suppression. Nodes receiving packets with this indicator set know that at least one preceding node explicitly requested 'no new-route-suppression'. This is a valuable HMP feature because it provides a safeguard against potential over-suppression of new route creation that may result in limited connectivity. The fourth parameter is the $\mathrm{NUM}_{\mathrm{ENOUGH\text{-}NEIGHBOR}}$ that prevents the HMP algorithm from being too aggressive.

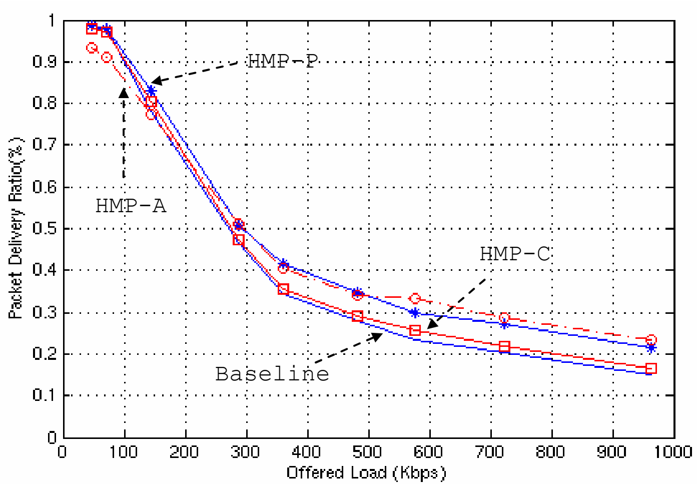

Figure 4-16: PDR Trace of HMP-A, HMP-C, HMP-P and Baseline System Compared

Figure 4-16 shows the PDR traces for the four different HMP systems under discussion. HMP-P and HMP-R are described in Section 4.5 while HMP-C and HMP-A represent HMP-Conservative and HMP-Aggressive HMP system configurations, respectively. HMP-A is literally an aggressive version of HMP-P that quickly determines hotspots (i.e., N_{THRESH} = 3, $M_{DELAY\text{-}THRESH}$ = 10 msec, $NUM_{ENOUGH\text{-}NEIGHBOR}$ = 4) and without utilization of the path_indicator. HMP-A is equally effective as HMP-P when the network is heavily loaded but results in limited connectivity when lightly loaded. As a consequence of its aggressiveness, HMP-A supports only 91 % PDR even when the offered load is only 72 Kbps. At the slightest indication of congestion, HMP-A suppresses route creation resulting in limited connectivity among mobile nodes.

HMP-C is a conservative version of HMP-P that utilizes the path_indicator and configures the HMP parameters as follows: N_{THRESH} = 8, $M_{DELAY\text{-}THRESH}$ = 100 msec, $NUM_{ENOUGH\text{-}NEIGHBOR}$ = 8. As shown in Figure 4-16, HMP-C closely resembles the baseline system and only shows slight improvements. If the HMP protocol is too aggressive it limits connectivity and if HMP is too conservative, it rarely detects hotspots and degrades to the baseline system performance.

As mentioned in Section 4.2.2 HMP relies on TCP throttling when severe congestion persists. As observed in Figure 4-17, HMP-TT (for HMP-TCP-Throttle) can selectively throttle TCP flows to relieve hotspots. TCP throttling is meaningful in the presence of congestion since TCP flows are typically transported as a best effort service where traffic rates are often transparent to higher layers (i.e., applications).

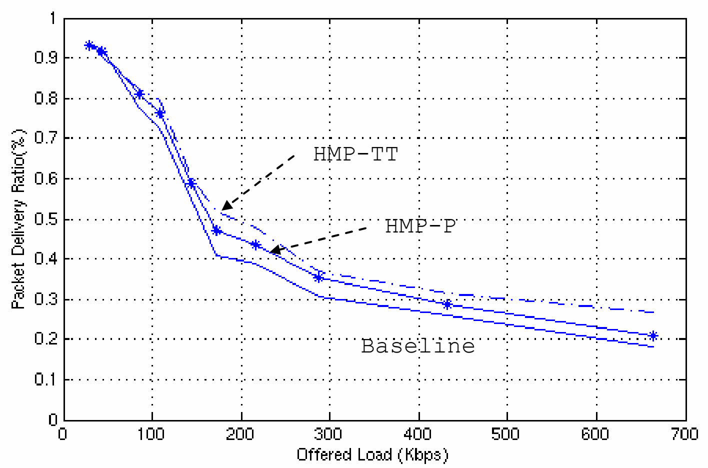

Figure 4-17: PDR of HMP-P, HMP-TT and Baseline System Compared

4.5.6 Proactive Routing Analysis

HMP is initially designed to operate with best effort on-demand routing protocols, as discussed in the previous sections. In what follows, we discuss the impact of HMP on the performance of two proactive routing protocols, namely, Destination Sequenced Distance Vector (DSDV) [36] and Optimized Link State Routing (OLSR) [37]. DSDV relies heavily on the exchange of periodic (and event-driven) HELLO messages, which essentially carry the routing table known to the sender. This exchanged routing information creates and updates routing state. Therefore, timely broadcast of HELLO messages is crucial to DSDV performance. However, HELLO messages result in substantial routing overhead. We use HMP to suppress these HELLO messages at hotspot nodes such that hotspots can temporarily avoid being used as a route. We use identical detection mechanisms to identify hotspots, as discussed previously.

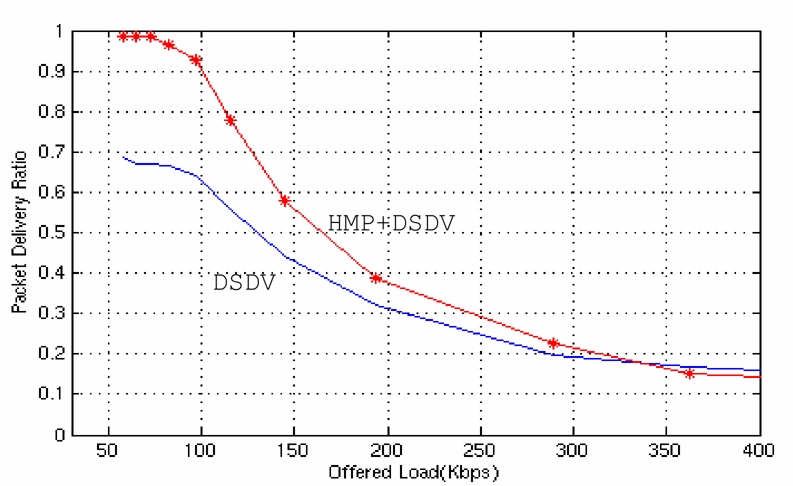

Figure 4-18: Packet Delivery Ratio Compared

The impact of HMP on DSDV performance is shown in Figure 4-18 and Figure 4-19. Figure 4-19(a) clearly shows that HMP+DSDV reduces the amount of control/routing overhead in terms of number of packets exchanged. However, when we observe overhead in terms of number of bytes, HMP+DSDV results in more overhead as shown in Figure 4-19(b). The reason for this anomaly is due to the size of HELLO messages in two systems. The average size of HELLO message in HMP system is substantially larger than that of baseline DSDV, indicating that the HMP system reveals more routing information than the baseline system. This is an important result because it indicates that HMP functionality effectively reduces the number of DSDV routing messages but each routing message conveys more comprehensive routing information.

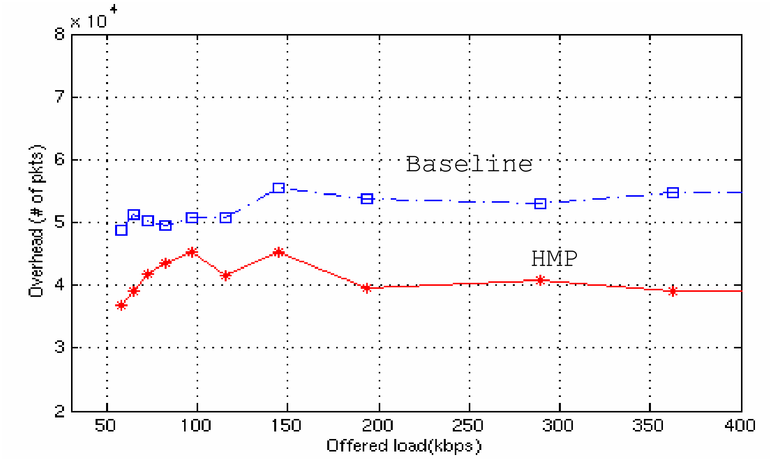

Figure 4-19(a): Routing Overhead (packets)

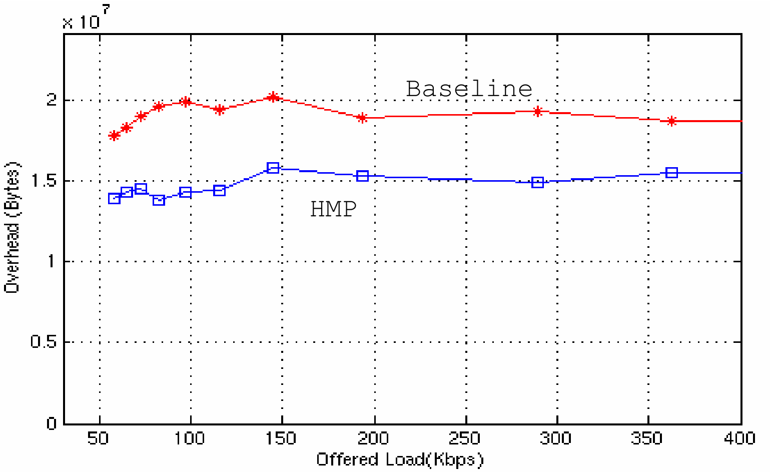

Figure 4-19(b): Routing Overhead (bytes)

Figure 4-20 shows the impact of HMP on OLSR, which is essentially an optimized version of pure link state routing where multipoint relays (MPRs) [37] are used to reduce the size of the information carried in a control messages and number of transmissions in the network. OLSR is more efficient than DSDV because in OLSR, all packets are routed only through MPRs, which results in reduced flooding. HMP can only be applied to OLSR through modifications to the OLSR algorithm. This is in contrast to the other protocols discussed in this chapter where HMP is implemented independently of the routing scheme.

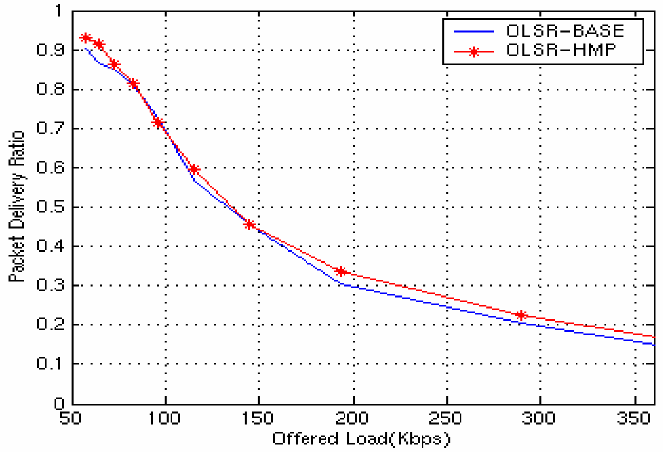

Figure 4-20(a): PDR comparison

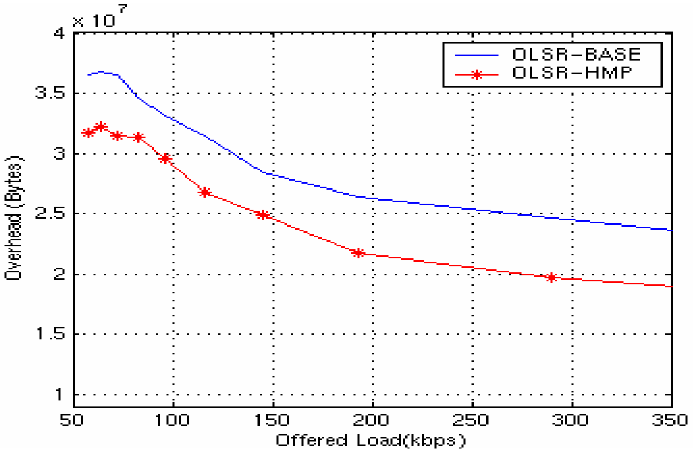

Figure 4-20(b): Overhead comparison

In contrast to DSDV, HMP does not have substantial impact on the performance of OLSR because MPRs are often identified as hotspots. If HMP declares an MPR a hotspot OLSR reacts to this by introducing additional MPRs to provide full connectivity. In such a case HMP also increases the overhead because MPRs generate routing messages to inform each other and update routing tables. Any increase in number of MPRs also indicates that more nodes are participating in the network-wide broadcasting of routing packets. Under such conditions HMP forces OLSR to select more MPRs, generate more routing messages and forward more routing messages in the system. HMP can only provide improvements to OLSR only when HMP effectively mitigate hotspots without introducing additional MPRs. This case is shown in Figure 4-20(a) and Figure 4-20(b). However, the trade-off between hotspot avoidance and creating more MPRs limits HMP's ability to provide constant improvements.

4.6 Experimental Wireless Testbed

To best understand the performance of HMP in a practical setting, and deployment issues with real wireless networks, protocols, and applications, we took a hands-on approach coupled with the analysis discussed in the previous section. In what follows, we discuss results from a small-scale wireless testbed implementation of HMP. The testbed consists of 8 notebooks running Linux (Red Hat 7.3 [60]) with Aironet wireless cards. We use the AODV v6.0 [63] released by Uppsala University and incorporate HMP modules discussed in Section 4.5. As part of our testbed methodology, we can emulate hotspot conditions at any nodes in the network simply by looping packets multiple times and introducing arbitrary delays with random packet loss. MACKILL [54] allows us to create various testbed topologies, as illustrated in Figure 4-21. Without such mechanisms, all nodes would be within transmission range, resulting in single-hop communications.

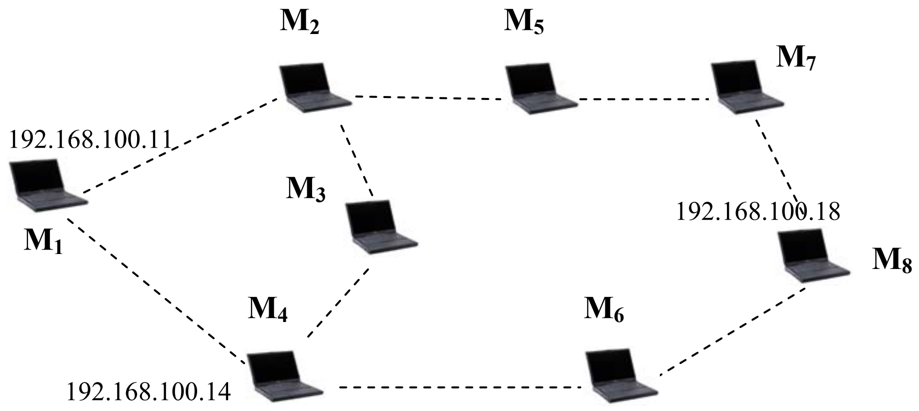

Figure 4-21: Testbed Topology

In Figure 4-21, node $M_{1,}$ and node M_8 represent the source and destination nodes while nodes $M_{4,}$ and M_6 represent the designated hotspots. We first traced the route between the source and destination nodes of the baseline AODV implementation to confirm that the system always routed though the shortest path (i.e., M_1-M_4-M_6-M_8) regardless of the existence of hotspots in the network. This is because AODV utilizes hop count as its route-creation metric. Under identical conditions we traced the route between the source and destination pairs for the HMP system and verified that HMP effectively aided AODV to take the non-hotspot path, (i.e., M_1-M_2-M_5-M_7-$M_{8.}$) rather than the shortest path. Under the same conditions we also created additional routes between M_3 and M_8 and observed that M_3-M_2-M_5-M_7-M_8 is always selected. Similar results are also observed when we introduce M_6 as a hotspot. For this scenario we collected measurements, as shown in Table 4-1. The table shows the average values for 5 identical runs of each testbed experiment, where each experiment encounters different interference levels with random delays and packet loss at the hotspots. We conducted ICMP tests by generating *ping* packets 1000 times between the source node M_1 and destination node M_8 at the rate of 4 pings per second. TCP tests comprised of downloading and playing of a 1 Mbyte MPEG movie from the source node (running Apache HTTP server) to the destination node. A UDP experiment comprised of transporting and playing the same MPEG movie but in UDP format. The table compares the baseline and HMP systems under identical conditions.

Traffic Type	System Type	Packet Loss	RTT(msec) Min/Average/Max	Routing Overhead (bytes)	Overhead Improvement (%)	PDR
ICMP	Baseline	132	35.2/760.1/4590.8	-	-	0.868
	HMP	0	5.2/7.9/430.7	-	-	1.000
UDP	Baseline	399	-	14672	28.16 %	0.564
	HMP	0	-	10540		1.000
TCP	Baseline	33	-	13664	25.75 %	0.956
	HMP	16	-	10146		0.982

Table 4-1: Testbed Results

It is observed from Table 4-1, that the ICMP test for the baseline system results in 132 packet losses out of 1000 packets transmitted, corresponding to a 13 % packet loss. For the baseline system, the minimum round trip time (RTT) measured is 35 msec, whereas the average round trip time is 760 msec. In contrast, for HMP system, minimum RTT is reduced to 5.2 msec and average RTT measurement is 7.9 msec with no packet loss. The UDP test also shows the benefit of HMP with no packet loss recorded, and with an overall reduction of 28 % in the routing protocol overhead. In contrast, the corresponding results for the baseline system result in 399 packet losses and packet delivery ratio of 56 %. Similar results are observed for the TCP test where HMP provides a 26 % reduction in routing overhead. However, packet delivery ratio is similar to that of HMP system because TCP packets are retransmitted when deemed lost.

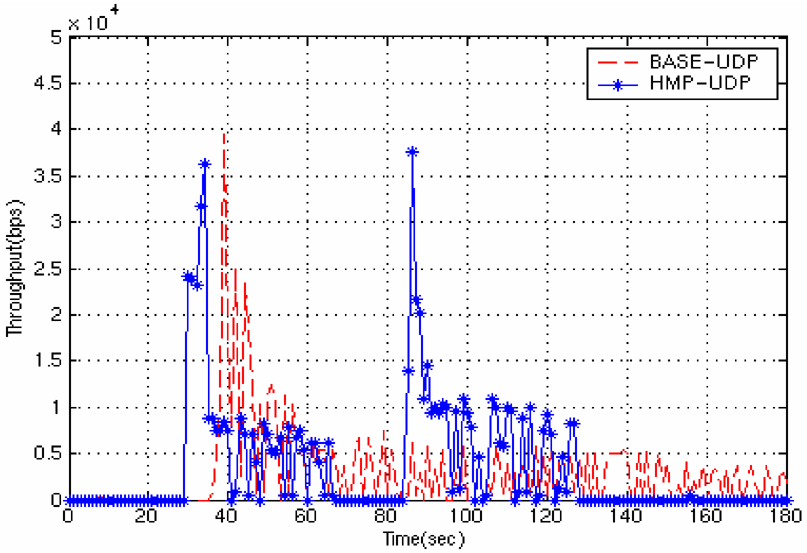

Figure 4-22(a): UDP Throughput Trace

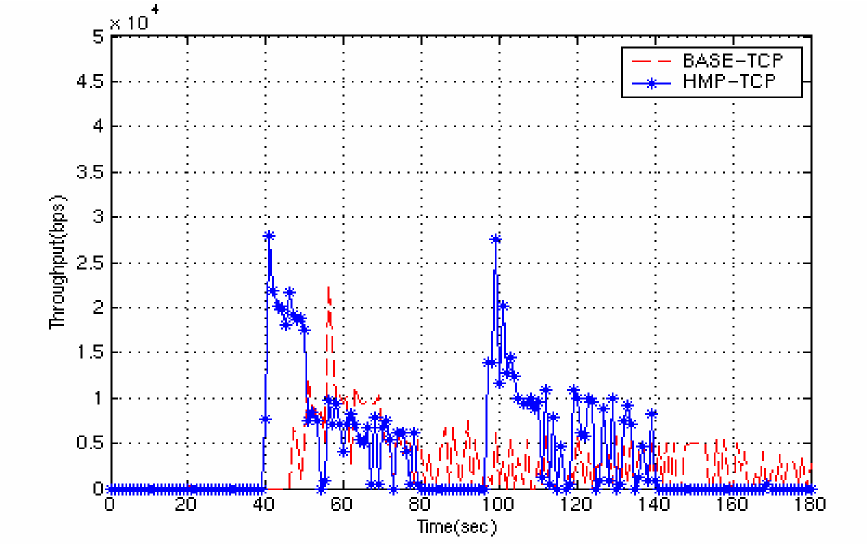

Figure 4-22(b): TCP Throughput Trace

Figure 4-22(a) shows the throughput traces of the streaming MPEG video experiment using UDP with and without HMP. Similarly, 4-22(b) shows the throughput traces for TCP download with and without HMP. The MPEG video download is repeated twice with 15 seconds of pause time in between downloads. Figures 4-23(a) and 4-23(b) indicate that HMP provides improvements over the baseline system in both cases (i.e., TCP and UDP tests). The throughput trace for UDP download using the baseline system takes a considerably longer time than the HMP system with the pause time being indiscernible in the trace. The poorer throughput is a result of the hotspot conditions encountered on en-route, which essentially causes excessive delay and increased packet loss. Due to the poor performance the video application could not playout the stream correctly. In contrast, the HMP avoided the hotspot nodes with the MPEG flow re-routed through an alternative path, (i.e., M_3-M_2-

M_5-M_7-M_8). In the baseline system the MPEG flow traversed the M_4 hotspot node. In contrast, the flow is routed around M_4 using HMP with fewer packet loss and delays observed. From Figure 4-22(a) and 4-23(b), the pause time between the two downloads is clearly discernible in the HMP system but not in baseline system.

4.7 Conclusion

In this chapter, we have presented a simple protocol that works with existing best effort routing protocols to mitigate hotspots in ad hoc networks. We have demonstrated through simulation that hotspots exist even in lightly loaded ad hoc networks and their existence can severely limit the performance of these networks. HMP tackles the problem of hotspots right at the point of congestion, as opposed to traditional end-to-end approaches found in the literature. We argued that traditional remedies such as end-to-end congestion control are often not effective in ad hoc networks and can limit the utilization and connectivity of the wireless network in the face of hotspots.

We evaluated HMP using both on-demand and proactive routing protocols. HMP provided significant increases in network performance and connectivity with lower routing overhead for ADOV and DSR. In the case of proactive routing schemes, HMP provided some performance boosts for DSDV but had limited success with the OLSR protocol due to its design of routing packets through specially designated nodes. To get some hands-on experience with the protocol we implemented HMP with AODV in a small-scale wireless testbed and confirmed the performance benefits observed under simulation. Based on our results, we recommend that future mobile ad hoc routing algorithms should incorporate the notion of hotspots directly into their protocols rather than simply adopting shortest path routing.

In the next chapter, we shift our research effort to the realm of sensor networks. Although a sensor network can be viewed as a part of the broader wireless ad hoc network family, its extreme deficiency in the resource availability (i.e., energy and bandwidth) makes it rather different from the mobile ad hoc networks. Among the many interesting issues, we investigate the reasons for poor information delivery (i.e., fidelity) in sensor networks, and propose a new routing protocol that enhances the information delivery.

Chapter 5

Solicitation-based Forwarding for Sensor Networks

5.1 Introduction

Recent technological advances in wireless communications make it possible for low cost, low complexity sensor networks to monitor and to detect environmental and tactical events. Sensor devices are typically equipped with a low power communication transceiver and a limited processor to facilitate signal processing. Because a sensor network can be deployed anywhere, even in areas where accessibility is limited, it is suitable for many emerging applications. One class of widely deployed applications is event-driven applications that are used to detect and report important events that occur in a sensor field. This type of application offers minimal traffic load and spends most of its time in an idle state. When an event is detected, the network becomes active and generates temporally and spatially correlated information that needs to be delivered to the sink. Since an event may be short-lived, the burst of information the network generates/senses during this time is likely to be of most importance to the application. A sensor network is therefore tasked to deliver a sufficient amount of information within a bounded time, i.e., fidelity [12]. However, numerous technical challenges hamper the delivery of adequate fidelity at the sink points in sensor networks. One of technical barriers to supporting sufficient fidelity comes from network dynamics. Network dynamics appear in various forms, e.g., wireless error, node failure, or anything that unexpectedly impedes on-going communications. Even when conducting indoor experiments, we often observe that only a fraction of the generated events are delivered to the sink due to the observed network dynamics. The presence of transitional regions [85], packet collisions, the funneling effect [88], and congestion [88] further limits the performance of sensor networks. A transitional region comprises highly unpredictable links with intermittent and asymmetric connectivity, which present significant networking problems. Sensor networks often exhibit non-isotropic radio ranges [86] and comprise asymmetric and unidirectional links. These conditions impair support of adequate levels of fidelity because link-layer reliability (or goodness of the link) is typically perceived through signaling exchanges or overhearing between participating nodes.

Adequate fidelity requires that event flows are routed through the "good-conditioned" nodes that form paths to the sink. The term good-conditioned may represent the energy-reserve of a sensor node, congestion status, routing distance, or any characteristic that correlates positively with the ability to deliver information to the sink. Sensor networks need cost-effective mechanisms to exploit these better-conditioned nodes to deliver information. Responsive self-configurability is another key property for fidelity support in sensor networks. A sensor network should be able to configure itself quickly and facilitate information delivery as soon as it is deployed. Moreover, a sensor network should be able to quickly respond to changes in network topology. A sensor network should also be responsive to nodes that fail over time which typically alter the connectivity graph of the network. Therefore, the delivery path needs to quickly reflect any observed changes in the topology and quickly adapt its delivery path to sustain event flows of information to the sink. Similarly, when new sensors are added to existing networks, they should be quickly integrated into the network with minimal overhead. Many of the existing routing protocols implemented in experimental sensor networks are not responsive to these challenges. Rather they incur a large control overhead, and lack the agility to cope with network and link dynamics (i.e., node failure, packet loss, link loss, new nodes, etc.). As a result this significantly impacts the fidelity of the delivered signal to the sink and sensor applications. To address these issues, we propose a new routing algorithm called *solicitation-based forwarding (SOFA)*. Through expensive Mica2 mote testbed experiments, we show that the on-demand nature of SOFA makes it cost effective, and responsive to network dynamics while supporting improved fidelity at the sink in comparison to existing experimental sensor network routing protocols [12] [15].

The structure of the chapter is as follows. Section 5.2 presents networking problems that motivate our proposal. The related work is presented in Section 5.3. This is followed by a detailed description of SOFA's operations in Section 5.4. Section 5.5 presents the experiemental evaluations of SOFA followed by concluding remarks.

5.2 Forwarding Problems in Sensor networks

In what follows, we discuss a number of forwarding problems found in experimental sensor networks, which motivate the design of SOFA. We use results from a set of experiments conducted on an experimental 36 Mica2 [39] mote testbed arranged in a dense 6x6 grid topology to quantitatively study forwarding problems in experimental sensor networks. The testbed software comprises the standard release of TinyOS [38], the Surge application, the MultiHopRouter [15] routing protocol, which is based on link quality estimation, and B-MAC [15]. Link quality estimation requires nodes to periodically broadcast beacon signals to create and manage per-neighbor statistical records of past communications that are used when evaluating link quality and making forwarding decisions at sensor nodes. Although these proactive approaches generally provide good routing paths for a stable network, they also present a number of limitations. First, they are cost-ineffective because they require all nodes to periodically exchange broadcast messages regardless of the level of network activity. Any transmission/reception consumes energy and bandwidth. The smaller the amount of sensor and control traffic in the network, the less energy consumed and probabilistically less collision observed. Therefore, it is important to keep the control overhead to a minimum in energy-limited sensor networks. Link quality estimation requires periodic signaling (or beaconing) and continuously consumes energy even when the network is in an idle state. This is counterintuitive because maintenance of unutilized paths only wastes energy.

Creating a forwarding path based on a statistical record may be time-consuming because a relaying node (i.e., parent) has to be determined at each wireless hop and these piecewise decisions take time to converge. Consequently, path convergence between a sensor and a sink often takes a substantial amount of time, preventing sensor devices from immediately reporting on-going phenomena after deployment. In our testbed, when using the default MultiHopRouter [15] routing protocol distributed with the TinyOS release the path convergence often requires several minutes and scores of routing message exchanges. Figure 5-1 presents an example of the path convergence distribution observed for 50 different experiments. Each experiment lasts for 30 minutes where we record the time to deliver the first packet to the sink, representing the path convergence time. A source node begins transmitting its data as soon as it powers on. The transmitted data from most source nodes is lost for an arbitrary period of time along partially constructed paths because the path to the sink is not completely resolved. Due to slow path convergence time, only 6% of source nodes achieve their path convergence within 60 seconds, approximately 50% in 120 seconds and the remainder spans up to the 10th minute. Such path convergence behavior exhibited in sensor networks poses a serious technical barrier for many applications because information delivery is preceded by a long settling time after network deployment or network dynamics, such as, link or node failure. Similar problems are observed when new sensors are added to an operational network. Typically, large convergence times are experienced when integrating sensors into a network. Similar path convergence issues occur when node failure occurs (e.g., energy depleted node) on a forwarding path, where the impact may last for a long period of time because it requires multiple samples to detect the loss of a next hop node and even more samples to acquire a replacement next hop node. During this time data packets may be continuously sent only to be lost. From our testbed results, we observe that the impact of node failure typically lasts for 3~5 minutes, and in the worst case the forwarding path never recovers (see Section 5.5.2 for a detailed discussion).

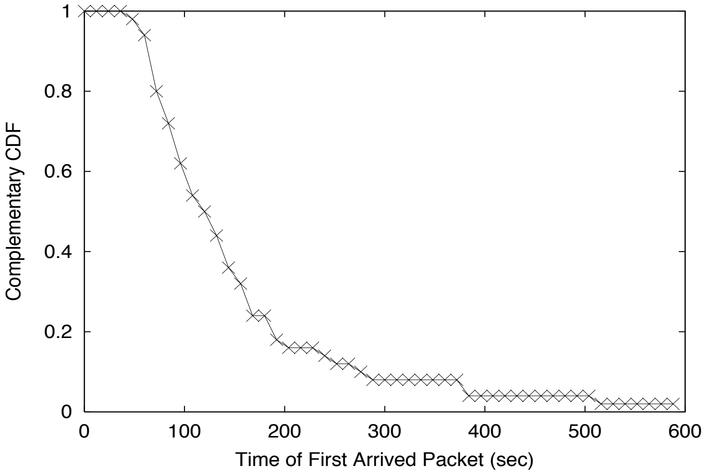

Figure 5-1 Path convergence of a proactive routing protocol

Another drawback of these proactive routing approaches is that they often fail to reflect link conditions at the exact time of the actual transmission. Events are rare in sensor networks and when an event occurs, a burst of information (i.e., an event train) is generated toward the sink node. However, estimation of link quality based on statistics from the recent exchange of periodic messages between nodes may not reflect the actual conditions when a burst of data traffic arrives at a link but is estimated when the burst of data is not present in the network. Therefore, it is likely that the link quality does not represent the actual condition when the data needs forwarding. We argue that forwarding decision should be made when the actual data is ready to traverse the wireless link. In other words, we argue that past measurements may have little relevance, particularly if they reflect past statistical states gathered under different conditions (e.g., idle state) from the actual data transmission.

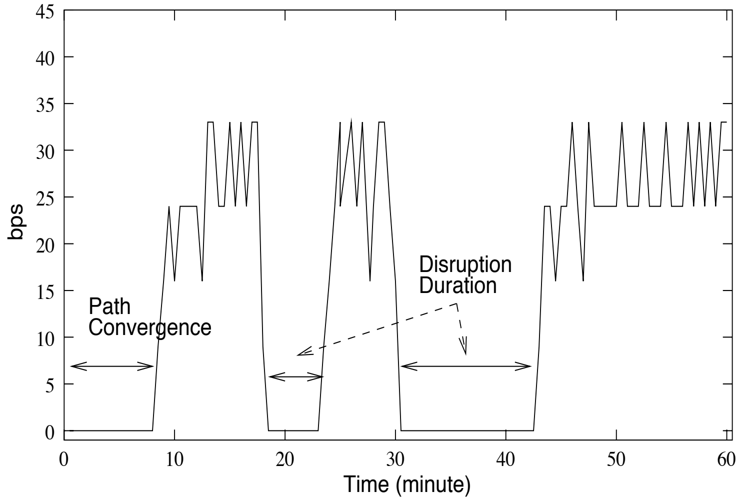

Figure 5-2: A monitored flow is traced at the sink to capture the impacts due to slow path convergence and node failures

The combination of path convergence and link quality estimate issues can substantially impact the overall performance of beacon-based proactive routing protocols [15]. Unless these forwarding problems are resolved, they limit the applicability of a sensor networks to a small number of simple low-fidelity applications (e.g., periodic reporting). Figure 5-2 shows a trace of a monitored event flow that encounters two route changes resulting from network dynamics (i.e., node failures in this example). As shown in Figure 5-2, the monitored event flow requires approximately 9 minutes for path convergence and the two re-routing conditions interrupt the event delivery for 4 and 12 minutes, respectively. Therefore, the event flow encounters aggregate disruption duration of approximately 25 minutes. This constitutes about 40% of a testbed runtime (i.e., 60 minutes). The main reason for such poor performance is associated with the link quality update interval. At low data rate, with intermittent collisions, nodes often do not resolve a valid relaying node, resulting in lengthy disruption of information delivery. This shortcoming can be somewhat improved if the frequency of routing message is increased but only at the cost of substantially increased control overhead.

5.3 Related Work

There are a number of routing protocols for sensor, mesh, MANET networks found in the literature. We first discuss the routing protocols released as part of the TinyOS software, and then discuss some relevant routing protocols for mesh and MANET networks. The TinyOS MultiHopRouter [15] protocol, which is widely used by the sensor network community, is based on the shortest-path algorithm that forms a spanning tree so that the path from any mote in the sensor field to the sink uses the least number of hops. Route control messages are periodically broadcast from each node in the network to estimate the routing cost and monitor link quality. The TinyOS Mintroute [15] protocol represents an adaptation of the simple MultiHopRouter protocol. MultiHopRouter uses the least number of hops as the primary metric with link quality as a tiebreaker, whereas, Mintroute uses the link quality with surrounding neighbors together with a cumulated route quality to the sink, ignoring the hop count in the route updates.

A number of researchers have revisited the design of routing protocols for mesh and sensor networks based on realistic wireless channel models founded on experimentation. In [89], the authors observe that the minimum hopcount without consideration of the channel characteristics shows poor performance. Destination sequenced distance vector (DSDV) [36] routing is a proactive routing algorithm that has influenced the implementation of TinyOS routing protocols. DSDV provides many-to-one routing to one destination at a time and can be used with either a hopcount metric or a quality metric. There are several geographical routing algorithms found in the literature that are applicable to sensor networks. Each node in greedy perimeter stateless routing (GPSR), for example, maintains a neighbor table that is updated via periodic beacon exchanges. However, these beacon messages constitute a large overhead for resource limited sensor networks. Other representative protocols in this class [93] [94] overcome the limitation of using periodic beaconing but still require some form of geographic coordinates provided by GPS for their operations. The geographic random forwarding protocol [95] represents one of the more sophisticated geographic routing solutions for sensor networks but its use of busy tones [96] makes it impractical to implement using standard sensor networking technology available today. There is a large body of work on MANET routing protocols [56] that has influenced our thinking, for example, the idea of "height" discussed in the next section is reminiscent of the TORA [32] routing protocol. However, these protocols are typically far too complex and costly in terms of control overhead to consider feasible for sensor network implementation.

5.4 SOFA Design

In what follows, we present the detailed design of the SOFA protocol.

5.4.1 Protocol Overview

SOFA establishes a path from a sensor to the sink based on hop-by-bop forwarding decisions by selecting appropriate relaying nodes at each wireless hop. A chain of relaying nodes composes the path to a sink. Each forwarding decision uses solicitation-based handshakes between a sender sensor and potential acceptors (i.e., next hop relaying nodes), where preference is given to the "best-conditioned" nodes as a relaying node at the time of packet communications. SOFA comprises four protocol phases; there are, *solicitation*, *acceptance*, *data-send*, and *passive acknowledgement*. In the solicitation phase, a solicitor seeks out a relaying node among its neighboring nodes by broadcasting a solicitation message called *solicit-to-forward (STF).* A neighboring node that is nearer to the sink receiving the STF accepts the solicitation by generating an *accept-to-forward (ATF)* message as long as it hears no other node has already responded to the STF. Once the solicitor node finds an acceptor, the accepting node becomes the *designated next hop (DNH)* for the solicitor node and solicitor node can send data to its DNH. Note, that the DNH is established on an on-demand basis and reflects the best link toward the sink at the time of data transfer across the link. In this sense the link is only assessed at the time of transmission and not continuously/periodically, which is the case for the link estimation schemes discussed in Section 5.2. The maintenance of the DNH is based on soft-state where the timer is associated with the event flow time-scale; that is, the DNH is kept active for a period of time to allow events to drain to the sink. After the soft-state timeout period the solicitor node would need to determine its DNH again. The thinking behind this is that the link quality may change after a certain period of time and the solicitor needs to determine its new DNH. This is triggered on an on-demand basis when the next event/data packet needs forwarding and is not assessed during the period when there is no data to transmit to the sink.

SOFA uses a passive acknowledgement mechanism, which means when a solicitor node overhears the forwarding of its data packet by its DNH it assumes reliable delivery has taken place. In the case it does not overhear the forwarding operation it can retransmit the original data packet if the application requires such a level of reliability. All transmitting nodes require a DNH and once a DNH is selected, data is unicast in a "distance-decreasing" direction toward a sink through a chain of DNH nodes. This is analogous to water flowing from higher to lower ground if we consider distance as height. We use the term *height* to represent the distance of a sensor node to the sink.

5.4.2 Height Initialization

During the height initialization phase, each node learns its height through sink-generated sink advertisement messages traversing the network. All messages (i.e., advertisements or any application query messages) that originated from a sink have a height of zero. As these messages propagate through the network their height information is incremented by one at each hop to reflect relative distance from the sink. Note, that height information of the sender is piggybacked in each message header. The sink node also has options to rebroadcast messages to update height information and re-advertise its existence. Any remote node or newly joined node that fails to receive advertisements would acquire its height through the height acquisition procedure described in Section 5.4.5.

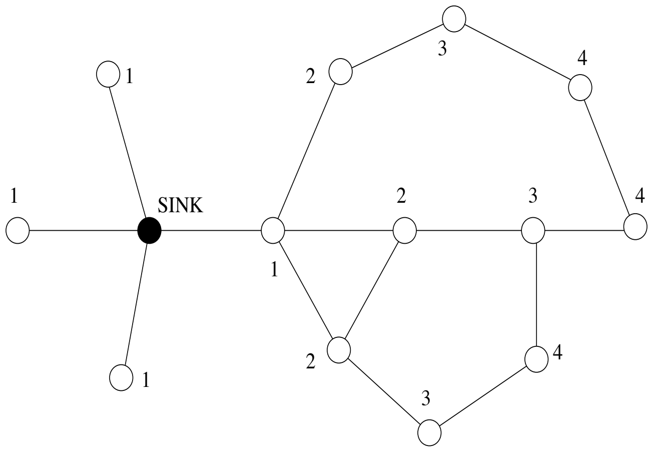

Figure 5-3: Height Initialization Result. Sink advertisement floods the network and each node learns their relative hopcounts (i.e., height) to the sink.

5.4.3 Solicitation Based Handshake

When a height-aware sensor node has data to transmit, it first checks whether it has a DNH. If it has a DNH, data can be immediately transmitted via its DNH, otherwise, the sender needs to acquire a DNH through solicitation-based handshake. A solicitation-based handshake starts when a DNH seeking node (i.e., solicitor) broadcasts an STF. When a node receives an STF, it first compares the height advertised in the STF with its local height and proceeds with an ATF response only if its height is less than that of the soliciting node (i.e., STF sender) to guarantee that the DNH is closer to the sink. In this context, neighboring nodes with small heights are termed next hop candidates.

Figure 5-4 shows an example of the solicitation-based handshake procedure. In Figure 5-4(a), node A has data to send but does not have a DNH, thus, node A broadcasts an STF to solicit a designated next hop. In this example, node B, node C, and node D are next hop candidates. As shown in Figure 5-4(b), node B is the first node to respond to the STF with an ATF. Note, that node D overhears the ATF response from node B and discard its pending ATF transmission. In other words, only one node within a common radio range responds to an STF. This mechanism provides a means to limit the number of ATF responses by permitting only a subset of next hop candidates to respond to single STF.

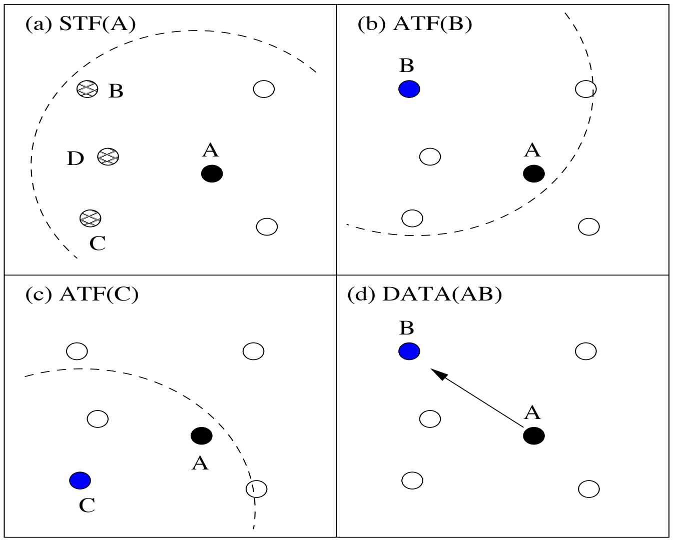

Figure 5-4: STF, ATF and DATA exchanges

In Figure 5-4, node C is outside the radio range of node B and unaware of the ATF response from node B. Therefore, node C also transmits an ATF to node A and as a consequence node A receives two ATF responses. In this example, node B is selected as the DNH of node A (DNH_A) because the ATF from node B reaches node A first. Having acquired a DNH, node A begins to unicast data toward the sink. The solicitation-based handshake completes when node A overhears node B (DNH_A) transmitting the data. Once the solicitation-based handshake completes, subsequent data messages from node A are unicast to node B without STF-ATF exchanges. Node A maintains its DNH unless an erroneous condition is detected. SOFA considers a number of packet losses, energy depletion, or any form of forwarding failure as erroneous conditions. SOFA conservatively assumes that a packet is lost when it is not passively acknowledged. For example, node A assumes that transmitted data is not received by DNH_A (i.e., node B of Figure 5-4) when it does not hear DNH_A relay its data packet. When γ consecutive passive acknowledgements fail, SOFA executes a new solicitation to acquire a new DNH. In such a case, the old DNH is blacklisted for a temporary period until a new DNH is selected. Note, that a blacklisted node can be reselected as a DNH if no other next hop candidates exist (i.e., no other ATF). If the solicitation process fails consecutively ϕ times, this may indicate the height of the soliciting node may be a local minimum (i.e., the node has no next hop candidates). This condition may arise when its DNH node expires or the network topology changes. SOFA resolves this situation through a height healing algorithm. Note, that both γ and ϕ can be tuned to make SOFA conservative or more aggressive. Details on height-healing and other height maintenance algorithms are discussed in Section 5.4.5.

5.4.3 ATF response and Defer-Time

One of the important features of SOFA is that ATF responses reflect current node conditions such that a DNH is less likely to be selected from problem-prone nodes. This mechanism is realized by coupling node conditions to the *defer-time* where defer-time is an additional waiting time that precedes an ATF response. Each next hop candidate receiving an STF delays its ATF response for the duration of its local defer-time. A node with a problematic condition would have a non-zero defer-time while a node in a better-condition would introduce no defer-time. This allows the STF sender to receive an ATF response from a better-conditioned node first. Equation (5.1) describes the defer-time of SOFA.

$$\text{defer\_time} = \{\text{random\% } CW_{SOFA}\}\cdot \text{slot\_time} \tag{5.1}$$

Note, that CW_{SOFA} (SOFA's Contention Window) can be used to represent various node conditions. For example, Equation (5.2) reflects a node's energy reserve status and congestion status.

$$CW_{SOFA} = (1\text{-}\beta)\cdot \text{energy}_{SLOT} + \beta\cdot \text{congestion}_{SLOT} \tag{5.2}$$

Note, that it requires congestion or energy concerns to have a nonzero $CW_{SOFA.}$ For example, when the energy reserves (i.e., $E_{CURRENT}$) of a node is below some predefined threshold value (i.e., E_{THRESH}), an additional energy_{SLOT} is added to CW_{SOFA}. These tunable system parameters are dependent on the hardware specifications and applications. Similarly, when congestion is detected at a node, the congestion_{SLOT} is added to CW_{SOFA}. The weight factor β is a system parameter to control the sensitivity of these metrics on the node's CW_{SOFA}. For example, when β = 1, CW_{SOFA} only reflects congestion condition whereas when β = 0, only energy.

Figure 5-5 captures the impact of the defer-time in DNH selection. We construct a 9-node network using the *ns-2* simulator [40] and observe DNH selection using (5.2). For simplicity, we set β = 0.5 and the congestion_{SLOT} to either 0 (i.e., no congestion) or 100 (i.e., congestion). The energy_{SLOT} is assigned with respect to a node's energy reserve. The simulation comprises one sink, one source node with height of 2, and seven intermediate nodes with height of 1. The next hop candidates (i.e., intermediate nodes) are assigned with diverse initial energy (i.e., node 7 = 10 J, node 3 = node 5 = 8

J, and the rest with 4 joules). The E_{THRESH} is set at 10 joules so that a non-zero defer-time always precedes an ATF response. Congestion is introduced randomly into the network in the first 50 seconds of the simulation run. With the source node transmitting 5 STF/second, we plot the corresponding DNH selections in Figure 5-5.

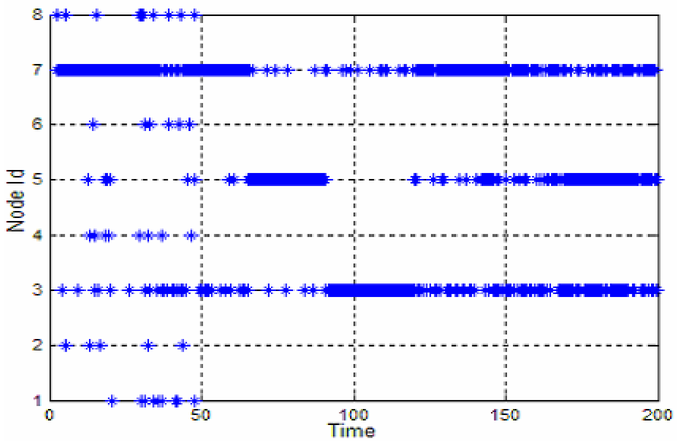

Figure 5-5: Impact of *defer-time* on DNH selections

The y-axis represents the node number of the selected DNH and the x-axis represents simulation time. Clearly, the three energy abundant nodes (i.e., node 3, 5, and 7) are predominantly selected as a DNH. Figure 5-5 also show that the existence of congestion in the first 50 seconds also has an impact on the DNH selection (i.e., DNHs are more distributed). In the first 70 seconds on the trace, node 7 is mostly selected as a DNH because it has the most energy but as energy reserve of node 7 decrease, node 5 and node 3 start to be selected as DNH. This result shows that the defer-time can effectively expose node conditions (e.g., energy, congestion, etc.) to the instantaneous forwarding decisions. Note, that in real testbed experiments discussed later in the chapter, only congestion condition is utilized because the energy-reserves are not accessible when using the Mica2 motes.

5.4.4 Height Maintenance

SOFA implements three simple height maintenance algorithms: height healing, height rollback, and height acquisition. Height healing resolves deadlock conditions, height rollback optimizes the height information, and height acquisition provides height information for a null-height node. In what follows, we discuss the algorithms.

5.4.4.1 Height Healing Algorithm

Height healing is executed when the height of a node becomes a local minimum and finds no next hop candidates. Such a node is dubbed a “sinkhole”. A sinkhole can become a DNH but never finds its own DNH. Although this condition rarely occurs, its impact is significant because all traversing packets are discarded at the sinkhole. Absence of a DNH would prompt re-solicitations but a sinkhole has no next hop candidates to acquire a DNH unless this anomaly is resolved. The only way to correct the condition is to increase the sinkholes height by one. This procedure is termed height healing. In general, height healing is preceded by multiple re-solicitation failures after a DNH is lost (i.e., unreachable). From our experimental results, most height healings is successful after one or two iterations.

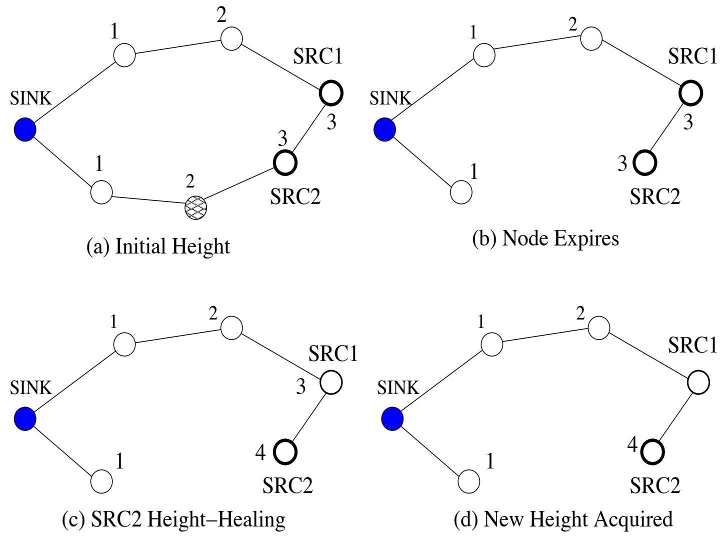

(a) Initial Height (b) Node Expires

(c) SRC2 Height–Healing (d) New Height Acquired

Figure 5-6: Height Healing Illustration

Figure 5-6 illustrates an example of the height healing algorithm. As illustrated in Figure 5-6(a), SRC1 and SRC2 are sending data to the sink. The initial heights of SRC1 and SRC2 are 3 in the example. As shown in Figure 5-6(b), SRC2 loses its DNH (e.g., due to node expiration, node failure) and data forwarding is suddenly disrupted. This condition is detected by SRC2 through the continuous lack of passive acknowledgements; therefore, the soliciting node assumes that its DNH is no longer reachable. SRC2 re-solicits for a new DNH but fails to acquire a DNH because its height is a local minimum. When the solicitation attempts continuously fail, SOFA identifies SRC2 as a sinkhole and performs height healing to resolve the anomaly. As illustrated in Figure 5-6(c), SRC2 increases its height by one and retries solicitation with the new height of 4. After a successful handshake of STF-ATF, SRC1 becomes the new DNH of SRC2 and data packets can now be forwarded toward the sink.

5.4.4.2 Height Acquisition Algorithm

Another important feature of SOFA is that it transparently integrates new nodes into existing operational networks. Newly joining sensor nodes do not have height information nor have knowledge of the sink. New nodes are considered to have null heights. There are two ways to acquire a valid height in SOFA. A new node can learn its height when it overhears any transmission from its height-aware neighbors. Its initial height becomes $\{h_{FIRST}+1\}$, where h_{FIRST} is the height of first overheard packet. The initial height is optimized through height maintenance algorithms. On the other hand, when a null height node wants to send data to the sink, it executes a height acquisition routine called “rippling”. Rippling involves transmission of an STF with null height (i.e., STF_{NULL}). When a height-aware node receives an STF_{NULL}, it generates a normal ATF with a height. Upon receiving ATFs from its neighboring nodes, the null-height node selects the minimum-height node as its DNH and sets its height to $h_{DNH}+1$ where h_{DNH} is the height of its DNH. However, when an STF_{NULL} is received by another null-height node (e.g., observed when cluster of new sensors are deployed) no ATF reply is sent; instead null-height nodes rebroadcast STF_{NULL} until it reaches a height-aware node. Each rippling node sets its ripple_flag and ATF response from the height-aware node backtracks to the STF_{NULL} originator through the path with the ripple_flag set. When the rippling phase completes, all associated null-height nodes becomes aware of their heights. Figure 5-7 illustrates the rippling case where node A, B and SRC2 represents newly joined nodes. SRC2 needs to report a detected event but lacks height information because it has not heard any transmissions from its height-aware neighbors. This condition triggers SRC2 to broadcast an STF_{NULL} and the broadcast message is received by node A and B. However, node A and node B cannot respond to the STF because they are also null-heighted nodes. In this case, node A and node B rebroadcast the STF_{NULL} (i.e., rippling). The process is repeated until an STF_{NULL} is received by a height-aware node. In this example, SRC1 and the sink respond to the rippled STF_{NULL} with ATFs. All ATFs piggyback the height information of the transmitter, as illustrated in Figure 5-7. Upon reception of ATFs, node A and node B learns their height. Since node A and node B have their ripple_flag set, they relay the ATF with their newly acquired height information. When SRC2 receives an ATF, the rippling routine completes and SRC2 acquires height.

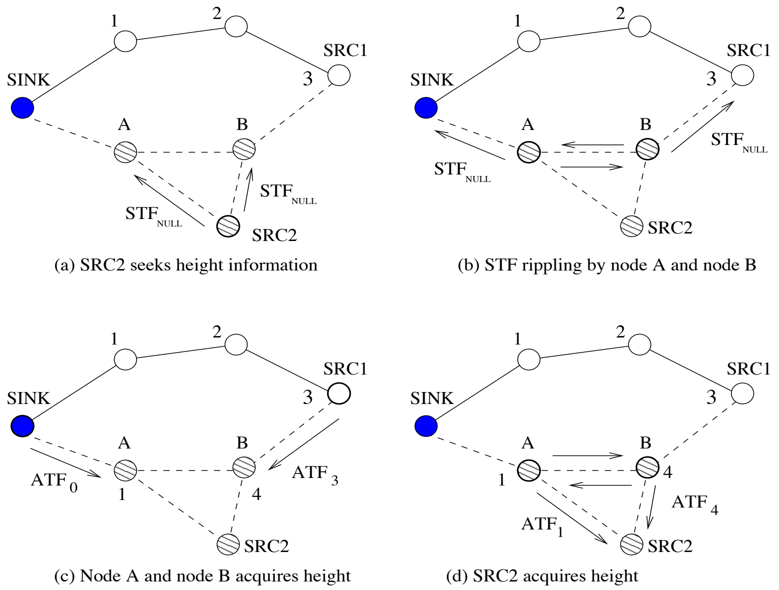

(a) SRC2 seeks height information

(b) STF rippling by node A and node B

(c) Node A and node B acquires height

(d) SRC2 acquires height

Figure 5-7: Height Acquisition Example

5.4.4.3 Height Rollback Algorithm

Height rollback is the counterpart of height healing algorithm that adjusts height to a smaller value. Height rollback is used to correct transient height sub-optimality in a sensor network. Typically, height sub-optimality is observed when new sensors are added into the existing network, altering the network topology. In such event, a node sometimes finds a new DNH with smaller height indicating that the new path is shorter in distance than the previous path. Soliciting nodes with height $h_{SOLICITOR}$ learns this condition when the height difference with its DNH node is greater than one (i.e., $h_{SOLICITOR} - h_{DNH} > 1$). This condition triggers the soliciting node to adjust its height to $\{h_{DNH} + 1\}$. From Figure 5-7, when node B seeks a DNH it sends out an STF (i.e., $h_B = 4$). When node B receives an ATF from node A, it detects that the height difference is more than 1 and adjusts its height to 2 (i.e., $h_B = h_A + 1$). Note, that height rollback does not incur any additional control load.

5.5 Experimental Testbed Evaluation

5.5.1 Mote Testbed Setup

In this section, we discuss the implementation of SOFA on a real sensor network using TinyOS [38] on Mica2 motes [39]. The testbed comprises 36 Mica2 motes arranged in a dense 6x6 grid topology. Two additional motes are strategically placed as snooper nodes to monitor communication details not captured by the sink node. Node spacing and transmission power are set such that one-hop neighbors can deliver more than 80%

transmitted packets, while two-hop neighbors deliver less than 20%. The data packet size is 36 bytes. We assume this experimental setting unless specified otherwise. We report detailed performance results of SOFA using B-MAC and compare it to MultiHopRouter [15] using B-MAC. MultiHopRouter is a routing protocol included in TinyOS for mote-based sensor networks where route control messages are periodically broadcasted from each node to estimate the routing cost and monitor link quality. We refer to the network running MultiHopRouter as the "baseline system" in the remainder of the chapter. In the testbed, we set γ (see Section 5.4.3) to be 3 and the solicitation limits to be 3 (i.e., ϕ = 3). However, the γ value changes to 7 when the link-layer retransmission option is enabled. These values reflect the link-layer retransmission limits [40] [56] and routing failure notification limits [40] [56] commonly adopted by the MANET [56] community.

5.5.2 Path Convergence Analysis

In this section, we compare path convergence of the SOFA system to that of the baseline system. We define path convergence of a flow as 'the time required for a packet to reach the sink for the first time'. In other words, the path convergence is, $t_{PATH\text{-}CONVER} = t_{FIRST\text{-}PKT} - t_{INIT}$, where t_{INIT} is the time when a source node generates a data packet for the first time and $t_{\mathrm{FIRST\text{-}PKT}}$ is the time when a packet from the source reaches the sink for the first time.

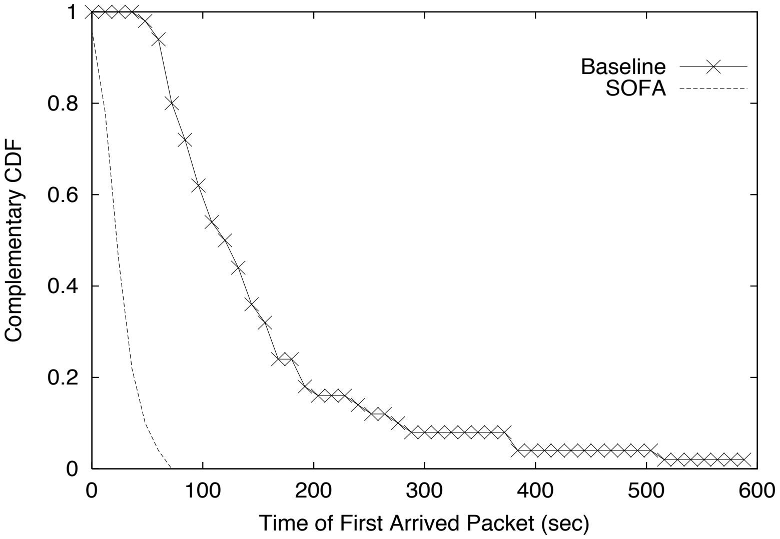

Figure 5-8: Path Convergence Time

Figure 5-8 shows the path convergence distribution of 50 different experimental cases in our 36-mote network. The x-axis of Figure 5-8 represents $t_{PATH\text{-}CONVER}$ and the y-axis represents the complementary CDF (cumulative distribution function). Figure 5-8 clearly shows that there is a significant difference in path convergences between the baseline and SOFA systems. In particular, 96% of SOFA's path convergences are accomplished within first 60 seconds while only 6% of path convergences are accomplished with the baseline system in the same window of time. In fact, the baseline system requires 372 seconds to achieve 96% of path convergences. The main reason for slow path convergence lies in its link quality update interval (i.e., periodic routing messages). At a low rate, with intermittent packet losses, nodes often fail to determine a valid relaying node. This can be somewhat improved if the routing update frequency is increase but only at the expense of substantially increased control overhead. For example, when the routing update frequency is doubled (i.e., update interval decreases from 20 seconds to 10 seconds), the average path convergence times improve by 37% for one of our baseline experiment but the corresponding control overhead increased by 200%. Moreover, increasing the update frequency may have a significant impact on fidelity because the increase in control load inevitably increases the collision probability and impairs the information delivery at the sink.

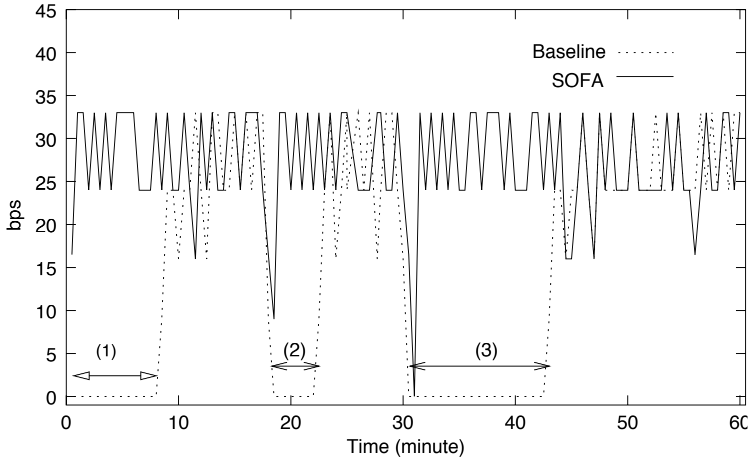

Figure 5-9: Comparison of a Monitored Flow

5.5.3 Network Dynamics Analysis

In what follows, we discuss the impact of network dynamics on information delivery at the sink (i.e., the fidelity). Since network dynamics cannot be controlled with Mica2 motes, we create artificial node failures that arbitrarily discard to-be-forwarded packets. In this set of experiments, only one source is present in the 36-node network. We carefully introduce node failures on the forwarding path at t = 18 minutes and t = 30 minutes, and observed the event flow for 60 minutes. The same sets of experiments are repeated for the baseline and SOFA systems. Figure 5-9 captures throughput of two systems, as measured at the sink. Note, that the baseline system takes approximately 9 minutes for path convergence (i.e., (1) in Figure 5-9). With the first artificial node failure at t = 18 minute, information delivery is interrupted for 4 minutes (i.e., (2) in Figure 5-9) until a new forwarding node is selected. Similarly, the disruption of information due to the second node failure at t = 30 minute lasted for approximately 12 minutes (i.e., (3) in Figure 5-9). The baseline system is slow in recovering a path because its link quality estimation mechanism requires multiple routing packets to realize and resolve the node failure problem. Note, that the baseline system uses the default settings for MultiHopRouter where each node broadcasts a non-propagating link-quality update message every 20 seconds. Thus, the link quality estimation of a particular link has an evaluation resolution of 20 seconds. This implies that path changes can only be executed in the multiples of 20 seconds (i.e., routing update rate). In contrast, the path convergence of SOFA complete in 10 seconds and the source node immediately starts its information delivery to the sink. More importantly, the impact of node failure on the SOFA system is minimal. When node failure is detected (i.e., loss of 3 consecutive passive acknowledgements), SOFA re-solicits acquires a new DNH in a single handshake that involved one STF and 2 ATFs. However, SOFA also incurs four additional re-solicitations before the second node failure and six more after second node failure. Lack of passive acknowledgements is misinterpreted as a node failure which triggers re-solicitation. This implies that SOFA entails additional control overhead without actual node failure when faced with packet loss. In fact, SOFA can entail substantial control overhead in a lossy environment. We discuss this issue later in Section 5.5.5.

When the network is in a stable condition with minimal problems, both systems perform well delivering about 80% of generated information. Under these conditions, SOFA's advantages over the baseline system are limited to faster path convergence, reduced overhead, and corresponding reduction in energy consumption due to reduced overhead. However, with the presence of any network dynamics, a greater disparity in performance begins to emerge. From Figure 5-9, the slow path convergence and two node failures constitute approximately 29% of total disconnected duration (i.e., no information is delivered during this period). In contrast, SOFA immediately achieves path convergence and the two node failures have virtually no impact on SOFA.

Figure 5-10 plots the packet reception ratio. Each point is an average of 10 experiments with 95% confidence interval. Note, that each testbed run lasts for one hour. The packet reception ratio (PRR) represents the ratio of number of packets received at the sink to the number of packets generated by source nodes. In this experiment, 6 nodes are randomly selected as source nodes. Their average height is 3 and their data rate is fixed at 1-packet/4-second. The x-axis represents the various degrees of network dynamics (as modeled in Section 5.5.3). The x-axis ranges from 0 NF (i.e., no node failures) to 9 NF (i.e., 9 node failures). Note, that node failures are only introduced on the 29 non-source nodes (i.e., not the 6 source nodes and sink node). We do not show results beyond 9 node failures since many of the experiments fail to deliver any data packets when there is more than 10 node failures, due to lack of connectivity in our testbed.

With zero node failures, SOFA shows a PRR gain of 7% over the baseline system. The main reason for this improvement lies in the path convergence where the baseline system has an average path convergence time of 8 minutes. With SOFA, all path convergences complete within 3 minutes. The impact of node failures on the two systems is clearly shown in Figure 5-10. As more network dynamics are introduced, the packet reception ratio of two systems degrades accordingly. However, SOFA provides improvements over the baseline system under all conditions. When there are 5 node failures, the baseline system can only support a PRR of 25% indicating that more than 4000 packets are lost in the network. In contrast, the corresponding PRR for SOFA is 40%, an improvement of 60% over the baseline system is achieved. In general, the SOFA system performs much better in the face of network dynamics. Average disruption duration for SOFA system is 28 seconds whereas the average disruption duration of the baseline system is 240 seconds.

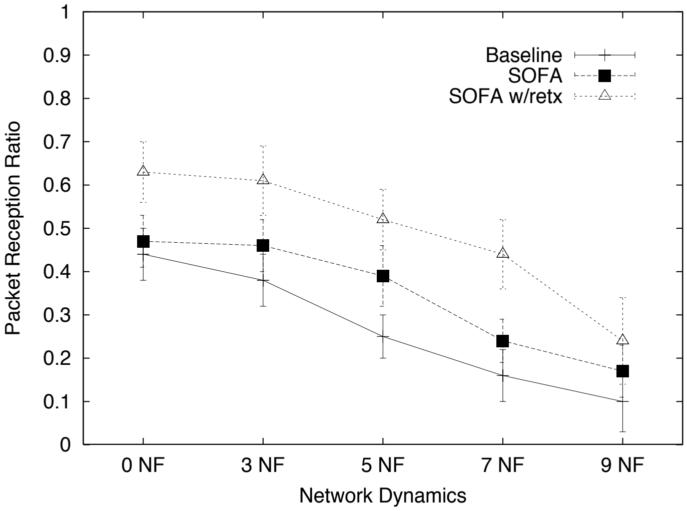

Figure 5-10: Comparison of Packet Reception Ratio

Figure 5-10 also shows the results of SOFA with link-layer retransmissions. As noted before, SOFA has an option to enable link-layer retransmissions. Consecutive retransmission failure triggers re-solicitation and when re-solicitations continuously fail SOFA-retx (i.e., SOFA with retransmission option enabled) assumes the node is a locally minimum in height and executes the height healing algorithm. As shown in Figure 5-10, the link-layer retransmissions for SOFA generally provide additional improvements in PRR. For example, with 3 NFs the PRR increases from 0.46 to 0.62, an improvement of more than 34 %. One interesting observation is that more than 60% of packet losses are observed in the vicinity of sink (i.e., nodes with h=1 and h=2) due to funneling effect [88] exhibited in sensor networks. We believe implementation of link-layer retransmissions in the vicinity of sink is a cost-effective way to improve overall fidelity.

5.5.4 Joining Node Analysis

To evaluate the integration of new nodes in the operation testbed, we conduct five sets of identical experiments on both systems. Each experiment starts with 24 active nodes with one source node. We let the network settle for 10 minutes (i.e., for path convergence of baseline system). In the 11th minute, we add a new source node every 2 minutes. In the 20th minute, we add cluster of 7 nodes (with one source node among them) simultaneously and we observed the integration behavior. Each experiment entails six integration instances using 12 sensor devices.

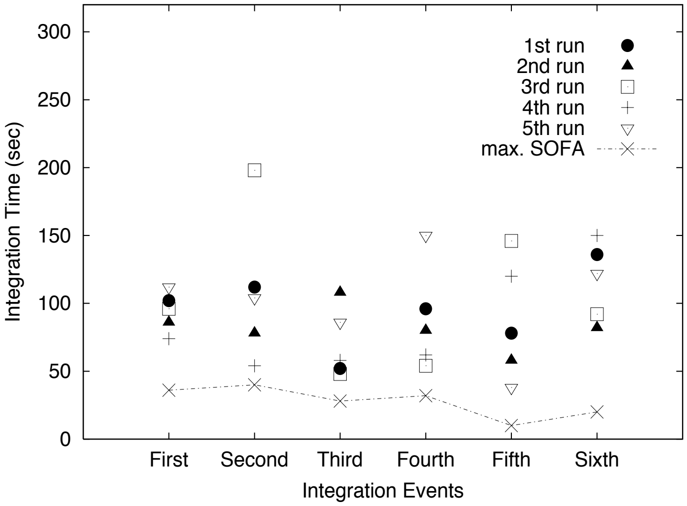

Figure 5-11: Integration Observation

As observed in Figure 5-11, the baseline system performs poorly and often requires long integration times. The average integration time for the baseline system is 94 seconds while the longest integration is 198 seconds. In contrast, SOFA completes all of its integrations within 60 seconds. Among the 30 integration opportunities, two rippling events are observed where STF traversed two null-height nodes for height acquisition. The integration experiments also entail 12 height rollbacks where interims heights are resolved during the DNH reacquisition processes. Figure 5-11 also plots the maximum integration times of SOFA.

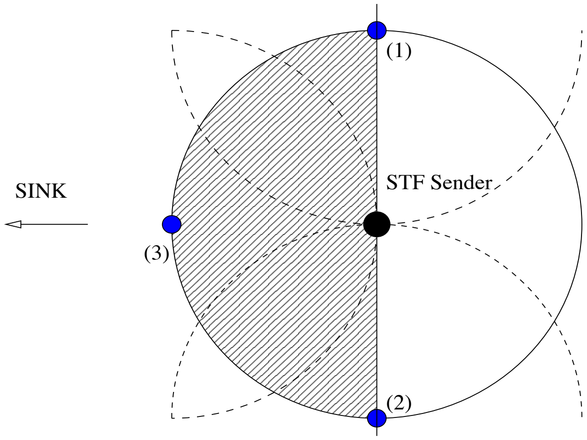

Figure 5-12: Possible positions of ATF senders when three ATF responses are sent (i.e., maximum bound). Note that next hop candidate can reside only in the shaded semicircle. Only one ATF can be generated in a common radio range.

5.5.5 Overhead Analysis

In this section, we discuss the overhead associated with baseline and SOFA systems. We consider all control and signaling messages to be overhead. Note, that the overhead of the baseline system has a constant rate because it broadcasts routing messages at a fixed interval. Therefore, the overhead of the baseline system is proportional to the network size (i.e., number of nodes) and operational duration. With the default settings of the baseline protocol, each node in the 36-node network generate routing message every 20 seconds. Therefore, the overhead of the baseline system has a constant rate of 6480 messages/hour, regardless of network activity. In contrast, the overhead of SOFA varies with the degree of activity in the network and is driven on an on-demand basis. When the network is in an idle state, SOFA does not produce any control overhead. Control overhead is associated only when an active sensor has data packets to forward. Active nodes are either source nodes or DNH nodes. In the 36-mote testbed, there are 6 source nodes with heights of {5, 4, 4, 3, 2, 2}. So, there are at most 20 nodes participating (because some node overlap) in solicitation-based handshakes. Each active node generates an STF message triggering a maximum of 3 ATF responses (see Figure 5-13). Thus, the control load for the SOFA system is at most 20×4=80 messages. As mentioned in Section 5.4, SOFA entails initial sink advertisements that flood the network. In the current SOFA implementation, the first sink advertisement broadcasts five consecutive advertisements at a 1-second interval. Therefore, with the worst-case assumption that all non-sink nodes rebroadcast advertisements without error, there would be at most 175 (i.e., 5×35nodes=175) additional control messages in the network. Even with the consideration of optional sink re-advertisements at 5 minute intervals, the final overhead for our experimental testbed is 80+175+(12×35) = 675 msgs/hour. This clearly outperforms the 6480 msgs/hour of the baseline system. The worst case for SOFA is when all nodes in the network are sources. Though we did not conduct experiments for this case, if all 35 non-sink nodes are sources then the number of STF/ATF control messages is 35 x 4 = 140, and the final overhead is

140+175+(12x35) = 735 messages/hour. Even in this worst case, the SOFA overhead is an order of magnitude lower than that of the baseline system.

Therefore, the SOFA system consumes significantly less control overhead but still offers PRR improvements, faster path convergence, and less service disruption. Reduction in control overhead correspondingly saves energy. In the event of forwarding failures, SOFA's overhead increases due to the re-solicitation process. Therefore, the more dynamic the network, the more control load SOFA generates. This behavior is shown in Figure 5-14 where the control load of both the baseline and SOFA systems is compared against the degree of network dynamics. The x-axis represents degrees of network dynamics in the form of node failures. The y-axis represents the corresponding control load produced for 60-minute testbed experiments. The control overhead of the baseline system is only dependent on the network size, regardless of the level of data traffic or packet loss. The control overhead for the baseline system shows a constant value of 6480 packets in all cases. In contrast, the overhead for SOFA varies with the number of sensed events and the degree of network dynamics. Under ideal conditions, SOFA generates approximately 700 control packets in the 60-minute testbed run, as described previously, but as observed in Figure 5-13 the resulting control overhead ranges from 1700 to 3300 packets. This is due to frequent loss of DNHs and lack of passive acknowledgment in the network. The control overhead of SOFA monotonically increases with network dynamics but the curve flatten out beyond 7 NF (i.e., 7 node failures). This is because the network contains less traffic due to lack of connectivity. Packets are simply not forwarded toward the sink and as a consequence less traffic exists in the network such that even with increase in network dynamics, fewer re-solicitations and height management procedures are triggered. We observe that the network often becomes disconnected when we introduce more than 10 node failures. In all cases, SOFA generates less control load than the baseline system.

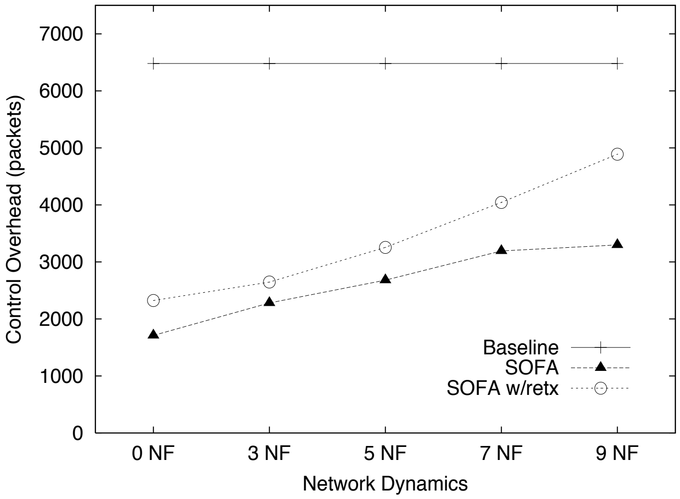

Figure 5-13: Overhead Comparison

5.5.6 Network Efficiency Analysis

In this section, we quantify the efficiency of the SOFA system through a parameter called the network efficiency. The network efficiency describes how efficiently a packet is delivered to the sink node with respect to total packet generation and it is defined by (5.3).

$$\eta_i = \frac{R^i_{SINK}}{(S^i_{SRC} + O^i_{NET})} \tag{5.3}$$

The parameter *i* of Eq. (3) denotes the type of the network (i.e., baseline or SOFA systems), R_{SINK} represents the total number of packet received by the sink, S_{SRC} denotes the total number of packets originated by source nodes, and O_{NET} represents the volume of control overhead in the network. Therefore, the upper bound for network efficiency is 1, which exists only if there is no packet loss and no control overhead in the network. If a network has substantial overhead and poor packet delivery, the network efficiency would be $\eta_i << 1$. Since a senor network inevitably entails packet loss and some control overhead, the network efficiency is typically well below the upper bound value of 1.

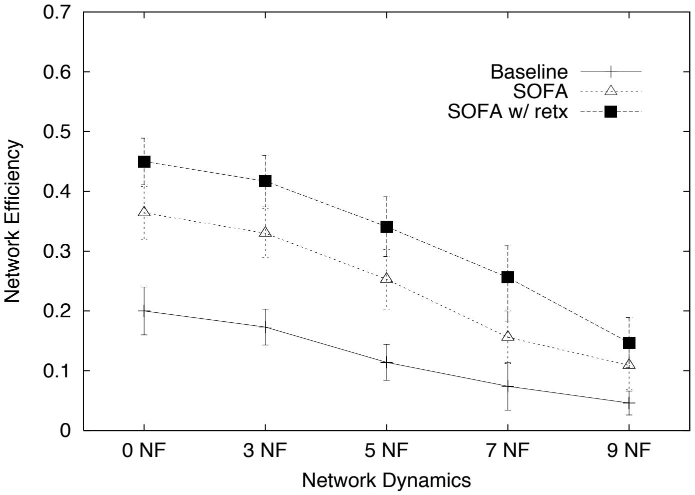

Figure 5-14: Impact of network dynamics on network efficiencies

Figure 5-14 plots the network efficiency of the baseline, SOFA, and SOFA-retx systems against network dynamics. Each point represents an average of ten experiments with 95% confidence interval. As observed in Figure 5-14, the efficiency of the baseline system (η_{BASE}) is 0.2 when there is no network dynamics present. As network dynamics increase the η_{BASE} begins to decrease correspondingly. For example, the average values of {R_{BASE} O_{BASE}, S_{BASE}} are recorded {2376, 6480, 5400} when no network dynamics are present but as the network dynamics increases to 5 NF, the R_{BASE} decreases to 1355 and consequently the η_{BASE} decreases to 0.11. The worst η_{BASE} of 0.046 is observed when the network dynamics is at 9 NF.

As shown in Figure 5-14, the SOFA system provides better network efficiency than the baseline system under all tested conditions. With no network dynamics, SOFA has η_{SOFA} of 0.364 while with 9NF η_{SOFA} decreases to 0.109. These results correspond to improvements of 82% and 137 % respectively when compared with the baseline system. The improvement over the baseline system is mainly due to the combination of control overhead reduction and R^{SINK} improvement. When the network faces little network dynamics, the dominant factor for the improvement is overhead reduction while with network dynamics present, the dominant factor is the R^{SINK} improvement. Figure 5-14 also plots the network efficiency of SOFA with retransmissions ($\eta_{SOFA-RETX}$). Enabling the link-layer retransmissions option increases the control overhead in all cases but also provides substantial improvement in R_{SINK}. In fact, the R_{SINK} improvement provided by the retransmission outweighs the impact of corresponding increase in control overhead. This condition is clearly shown in Figure 5-14 because $\eta_{SOFA-RETX}$ always outperforms η_{SOFA}.

5.5.7 Energy Analysis

In this section we observe how the energy is spent in a SOFA network. In particular, we are interested in how energy is dissipated in the network in support of the resulting information delivery under various experimental conditions. To measure the energy dissipation, we introduce a performance metric called *energy expense*. The energy expense represents the ratio of total dropped packets in the network to the total delivered packets. In other words, the energy expense describes the average energy consumption per delivered packets. Note that energy expense extends the energy tax metric used in [88]. The energy expense incorporates both data packets and control packets. Since packet transmission and reception consumes the most portion of the energy, the energy expense metric is a good indication of how the energy has been consumed.

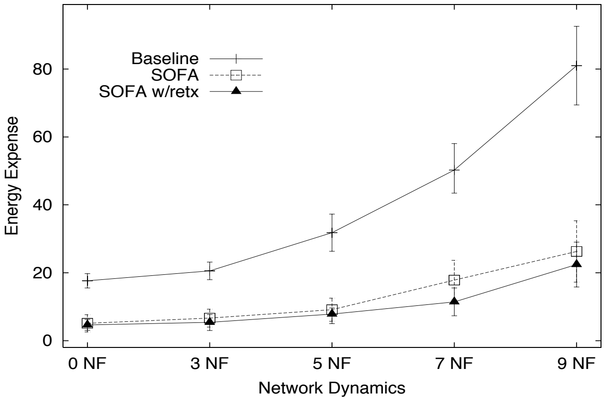

Figure 5-15: The energy expense describes how energy has been spent to support the resulting information delivery perceived at the sink.

Figure 5-15 plots the energy expense against the network dynamics. It is shown that the increase in the network dynamics makes the information delivery more expensive. With more network dynamics present in the network, more control overhead and more packet losses are encountered. For example, when the network encounters little network dynamics (e.g., 3 node failures) the average energy expense for the baseline network is approximately 21, indicating that 21 packets (i.e., control packet and data packets) are consumed per one delivered data packet. In contrast, in the SOFA network, only an average of 6.5 packets is required for delivering a data packet to the sink. This indicates that the SOFA network requires only 1/3 of energy consumption per delivered packet. The difference in energy expense widens further as the network dynamics increases as shown in Figure 5-15. The obvious reason is due to the difference in R_{SINK} and the number of dropped packets. Specifically, with 7 node failures, the baseline network entails energy expense of approximately 50 whereas with SOFA, the energy expense is recorded at 18. SOFA improves the energy expense by 64 %. In addition, when the retransmission is implementation, the SOFA network achieves energy expense of 11.4, improving the energy expense by more than 77 % when compared with the baseline network. Clearly, the SOFA networks (i.e., SOFA and SOFA-retx) increase R_{SINK}, increase the network efficiency, and improve the energy expense.

5.8 Conclusion

The topology of sensor networks continuously change and thus information delivery has to efficiently adapt to these changes while sustaining on-going communications with low overhead. In this chapter, we presented the design, implementation, and evaluation of SOFA, an on-demand solicitation-based forwarding protocol for sensor networks. SOFA represents a very simple and scalable solution for routing in experimental sensor networks. Our experimental testbed results confirm that SOFA provides excellent path convergence times and is responsive to various network dynamics experienced in sensor networks. We show through extensive experimentation that on-demand approaches such as SOFA are very applicable to event-driven sensor applications, and that SOFA outperforms the commonly used link estimation-based routing schemes implemented in TinyOS sensor networks.

Chapter 6

Conclusion

This thesis has covered three important issues in the broader area of wireless ad hoc networks that comprises mobile ad hoc networks and sensor networks. As of today, wireless ad hoc networks are gradually deployed in the real world and slowly emerging as a part of our daily lives. To cope with the unpredictable nature of our dynamically changing physical environment, wireless ad hoc networks should be able to adapt to changes in resource availability and overcome any unanticipated networking problems whilst satisfying a wide range of application requirements. Fulfilling these requirements in such environment is very challenging because the traffic present in the wireless ad hoc network is continuously affected by performance degrading network dynamics.

Chapter 2 presented our first contribution of this dissertation, the INSIGNIA QOS framework. To resolve the challenges and problems for QOS support in mobile ad hoc networks, we have investigated the solution space and proposed a QOS framework.

Given the fact that QOS guarantees are not feasible in mobile ad hoc networks, the INSIGNIA QOS framework is designed to support the adaptive service paradigm. The key component of the QOS framework is the *INSIGNIA signaling system*, an in-band signaling system specifically designed to address the QOS-related challenges in mobile ad hoc networks. The INSIGNIA signaling system has been widely recognized as the first contribution to the issues of QOS in mobile ad hoc networks and over the years, the INSIGNIA signaling system has become the benchmark signaling protocol to outperform in MANET community.

Several important contributions from the INSIGNIA QOS framework have influenced the QOS research in MANET community. First of all, the “in-band” signaling approach was first introduced by INSIGNIA QOS framework. It was shown that the in-band signaling mechanism is well-suited for supporting adaptive QOS in mobile ad hoc networks. Secondly, the “soft-state” resource management scheme was promoted by the INSIGNIA QOS framework. It was shown that the soft-state approach provided better network utilization whilst efficiently resolving the false restoration and resource lock-up problems in MANET. The INSIGNIA signaling system has successfully fused the in-band signaling, soft-state resource management, and per-flow state management to orchestrate fast reservation, fast restoration, and end-to-end service adaptation.

Chapter 3 presented a detailed performance evaluation of the IEEE 802.11 [50] based INSIGNIA signaling system with a number of MANET routing protocols. Extensive simulation studies using NS-2 simulator package [40] and hands-on experience from our experimental testbed have confirmed that INSIGNIA is suitable for supporting QOS in mobile ad hoc networks. It was shown that INSIGNIA provided operational transparency to a number of MANET routing protocols (e.g., AODV [30], DSR [31] and TORA [32]) and provided significant performance gains for various TCP (i.e., TCP-Reno [33], TCP-Vegas [33], and TCP-SACK [34]) and UDP flows. The simulation codes (i.e., NS-2 codes) and the testbed codes (i.e., INSIGNIA on DSR/FreeBSD and INSIGNIA on AODV/Linux) used for the study reported in this dissertation have been publicly available at the INSIGNIA project website [71] (www.comet.columbia.edu/insignia) for many years and they have been used by many MANET researchers.

In Chapter 4, we studied the congestion condition called a *hotspot*. A hotspot is defined as a node experiencing flash congestion conditions or a period of excessive contention conditions in wireless ad hoc networks. It was shown that hotspots exist even in lightly loaded mobile ad hoc networks and their existences severely degraded the network's performance. A thorough investigation has revealed that the existence of a hotspot is largely due to mobility in mobile ad hoc networks. The mobility of nodes continuously changed the network topology and caused the on-going traffic to reroute. This caused variations in network the loading conditions and caused transient congestion conditions called hotspots. The node mobility creates, removes, and even migrates hotspots in mobile ad hoc networks. These hotspots have caused packet loss, delay-spikes, and even triggered route maintenance (i.e., when misinterpreted as routing failure). As a solution to the hotspot condition, this thesis proposed the Hotspot Mitigation Protocol (HMP) that works with existing best effort routing protocols.

To best of our knowledge, HMP represents the first generic solution to hotspot problems in mobile ad hoc networks. HMP was the first protocol that formally identified the existence of hotspot conditions and it was the first protocol specifically designed to mitigate the hotspot conditions. Three integral parts of the hotspot mitigation protocol are; (1) accurate and cost-effective hotspot detection, (2) hotspot mitigation mechanism, and (3) a traffic throttling scheme. The hotspot detection mechanisms presented in this dissertation are based on combination of several congestion indicators (i.e., packet loss pattern, MAC-delay, and buffer occupancy) found in the network. The HMP effectively suppressed and dispersed new/rerouted flows from hotspot regions to mitigate the congested condition. In addition, HMP provides a traffic throttling scheme that rate controls the best effort TCP flows to relieve congestion condition. We have conducted a thorough evaluation of HMP with several MANET routing protocols and it was confirmed that HMP provided significant improvements in network performance (i.e., throughput, packet loss, delay, etc.), balanced resource usage, and reduced routing overhead. Based on our results, it is recommended that future mobile ad hoc network protocols should incorporate the notion of hotspots in their design considerations.

In Chapter 5, we shift our research focus to sensor networks, the foremost frontiers of wireless ad hoc networks as of today. Based on the observation that current routing algorithms for sensor networks provided poor information delivery (i.e., low fidelity), we have conducted a thorough investigation to identify the problem. In depth investigation using a TinyOS [38]/Mica2 [39] testbed has revealed that the poor fidelity is largely due to the unresponsive nature of route selection commonly practiced in sensor networks. To counter this problem, we proposed SOFA, an agile, cost-effective, and high-fidelity yielding hop-by-hop routing protocol that creates a path through a series of solicitation based handshakes that considers local conditions at each forwarding node.

Experimental testbed results confirmed that SOFA achieved fast path convergence at deployment and quickly acquired an alternative path with minimal signaling overhead when faced with path changing conditions. It was also shown that the path maintenance in SOFA is minimal and integration of new sensors completed immediately and seamlessly even with large number of new sensors. The combination of these attributes facilitated SOFA to provide much enhanced fidelity in comparison to the baseline network. Based on these results, we recommend that components in sensor networks should be made more agile. We also argue that adaptability and agility are the fundamental building blocks of routing algorithm for sensor networks in addition to the energy efficiency attribute. SOFA also provided significant reduction in energy consumption where the energy savings in SOFA network primarily came from decrease in signaling overhead. By eliminating the periodic beaconing, SOFA avoided a significant amount of periodic beacon messages and in turn, conserved a significant amount of energy. The on-demand nature makes SOFA cost effective; its agile self-adapting nature makes it resilient to network vagaries; and its use of timely solicitation-based handshakes make its forwarding decisions effective in data delivery.

Despite the significant advances in the broader area of wireless ad hoc networks, the issues addressed in this thesis are still under active research. It is envisioned that the advances in modulation techniques, antenna technology, hardware (i.e., MEMS), and better protocol design (i.e., routing, MAC, etc.) will allow the wireless ad hoc network to be deeply integrated into our daily lives. In this thesis, we have addressed a number of challenging but important issues for supporting adaptive QOS in wireless ad hoc networks. While much additional work remains to be done, this thesis offers important contributions to the vision of realizing tomorrow's system.

Chapter 7

My Publications as a Ph.D Candidate

7.1 Journal Papers

- Seoung-Bum Lee and Andrew T. Campbell, HMP: Hotspot Mitigation Protocol for Mobile Ad hoc networks, *Ad Hoc Networks Journal (ELSEVIER)*, Vol.1 No.1, 2003

- S.B. Lee, G.S. Ahn, X. Zhang, and A.T. Campbell, INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks, *Journal of Parallel and Distributed Computing (Academic Press)*, *Special issue on Wireless and Mobile Computing and Communications*, Vol. 60 No. 4 page 374-406, April 2000

7.2 Magazine Papers

- S.B. Lee, G.S. Ahn, A.T. Campbell, Improving UDP and TCP Performance in Mobile Ad Hoc Networks with INSIGNIA, *IEEE Communication Magazine*, June 2001

7.3 Conference and Workshop Papers

- Seoung-Bum Lee, Kyung Joon Kwak, and Andrew T. Campbell, Solicitation-based Forwarding for Sensor Networks, *IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON 2006)*, Reston, VA, Sept 2006

- Seoung-Bum Lee and Andrew T. Campbell, HMP: Hotspot Mitigation Protocol for Mobile Ad Hoc Networks, *International Workshop on Quality of Service (IWQOS 2003)*, Monterey, CA, June 2003

- Lee, S.B., Ahn, G.S., X. Zhang, and Campbell A.T., Evaluation of the INSIGNIA Signaling System, *In Proc. of 8th IFIP International Conference on High Performance Networking (Networking 2000)*, Paris, France, May, 2000.

- Seoung-Bum Lee and Andrew T. Campbell, QoS Issues in Mobile Ad Hoc Networks, *Seminar on Mobile Multimedia Communication - Systems and Networks*, Dagstuhl, Germany, May 1999.

- Seoung-Bum Lee and Andrew T. Campbell, INSIGNIA: In-band signaling support for QOS in mobile ad hoc networks, *In Proc. of 5th International Workshop on Mobile Multimedia Communications (MoMuC,98)*, Berlin, Germany, October 1998.

7.4 IETF Internet Drafts

- Seoung-Bum Lee, Jiyoung Cho, and Andrew T. Campbell, Hotspot Mitigation Protocol(HMP), Internet Draft, *draft-ietf-lee-hmp-00.txt*, IETF MANET Working Group Document, October 2003

- Seoung-Bum Lee, Gahng-Seop Ahn, Xiaowei Zhang and Andrew T. Campbell, INSIGNIA, Internet Draft, *draft-ietf-manet-insignia-01.txt*, IETF MANET Working Group Document, November 1999.

- Seoung-Bum Lee and Andrew T. Campbell, INSIGNIA, Internet Draft, *draft-ietf-manet-insignia-00.txt*, IETF MANET Working Group Document, November 1998

- Seoung-Bum Lee and Andrew T. Campbell, INSIGNIA, Internet Draft, *draft-lee-insignia-00.txt*, November 1998

References

[1] J. Freebersyser and B. Leiner, A DoD Perspective on Mobile Ad Hoc Networks, *Ad Hoc Networking*, C. E. Perkins, Addison-Wesley, 2001, pp. 29-51

[2] Crawley E, Nair R, Rajagopalan B, Sandrick H. A Framework for QoS Based Routing in the Internet. *RFC 2386*, August 1998.

[3] Braden R, Zhang L, Berson S, Herzog S, Jamin S. Resource reSerVation Protocol (RSVP) Version 1 Functional Specification. *RFC 2205*, September 1997.

[4] P. Karn, MACA: A New Channel Access Method for Packet Radio, *In Proc. ARRL/CRRL Amateur Radio Ninth Computer Networking Conference*, pp. 134-140, April, 1990.

[5] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, MACAW: A Media Access Protocol for Wireless LANs, *In Proc. of ACM Sigcomm'94*, pp.212-225, 1994.

[6] Fullmer and J.J. Garcia-Luna-Aceves, Floor Acquisition Multiple Access (FAMA) for Packet-radio Networks, *In Proc. ACM Sigcomm'95*, Cambridge, MA, pp. 262-273, 1995.

[7] Talucci and M. Gerla, MACA-BI (MACA By Invitation). A Wireless MAC Protocol for High Speed Ad hoc Networking, *In Proc. of IEEE ICUPC'97*, 1997.

[8] Andrew Muir and J.J. Garcia-Luna-Aceves, An Efficient Packet-sensing MAC Protocol for Wireless Networks, *ACM Jounal on Mobile Networks and Applications*, Vol.3, No.2, pp. 221-234, August 1998.

[9] J.L. Sobrinho and A. S. Krishnakumar, Quality-of-Service in Ad Hoc Carrier Sense Multiple Access Wireless Networks, *IEEE Journal on Special Areas in Communications*,Vol. 17, No. 8, August 1999.

[10] Sung-Ju Lee and Mario Gerla, Dynamic Load-Aware Routing in Ad hoc Networks, *In Proc. of 3rd IEEE Symposium on Application-Specific Systems and Software Engineering Technology*, 2000.

[11] C. E. Nishimura and D. M. Conlon, IUSS dual use: Monitoring whales and earthquakes using SOSUS, Mar. Tech nol. Soc. Journal, vol. 27, no. 4, pp. 13–21, 1994.

[12] Chalermek Intanagonwiwat, Ramesh Govindan, Deborah Estrin, John Heidemann, and Fabio Silva, Directed diffusion for wireless sensor networking, *ACM/IEEE Transactions on Networking*, February 2002.

[13] W.R. Heinzelman, J. Kulik, and H. Balakrishnan, Adaptive protocols for information dissemination in wireless sensor networks, *In Proc. of the 5th Annual International Conference on Mobile Computing and Networking*, pp 174-185, 1999.

[14] John Heidemann, Fabio Silva, and Deborah Estrin, Matching data dissemination algorithms to application requirements, *In Proc. of the ACM SenSys Conference*, pp. 218-229, Los Angeles, California, USA, November 2003. ACM.

[15] A. Woo and D. Culler, Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. *In Proc. of the 1st ACM Conf. on Embedded Networked Sensor Systems*, Los Angeles, CA, Nov 2003

[16] Benjie Chen, Kyle Jamieson, and Hari Balakrishnan, An energy efficient coordination algorithm for topology maintenance in ad hoc wireless networks, *In Proc. of the 7th Annual International Conference on Mobile Computing and Networking (Mobicom 2001)*, pp. 85-96, July 2001.

[17] Y. Xu, J. Heideman, and D. Estrin, Geography-informed energy conservation for ad hoc routing, *In Proc. of the 7th Annual International Conference on Mobile Computing and Networking (Mobicom 2001)*, pp. 70-84, July 2001.

[18] Grime S and Durrant-Whyte H F. Data fusion in decentralized sensor networks. Control Engineering Practice, 2(5): pp. 849-863, 1994.

[19] D. Guo and X. Wang. Dynamic sensor collaboration via sequential monte carlo. *IEEE Journal on Selected Areas in Communications*, 22(6): pp.1037-1047, August 2004.

[20] W. Ye, J. Heidemann, and D. Estrin, An energy efficient MAC protocol for wireless sensor networks, *In Proc. of the 21st International Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2002)*, pp. 1567-1576. New York, June 2002.

[21] Joe Polastre, Jason Hill, and David Culler, Versatile low power media access for wireless sensor networks, *In Proc. of Second ACM Conference on Embedded Networked Sensor Systems (SenSys 2004)*, pp. 95-107. Baltimore, November, 2004.

[22] T. V. Dam and K. Langendoen, An adaptive energy-efficient MAC protocol for wireless sensor networks, *In Proc. of First ACM Conference on Embedded Networked Sensor Systems (SenSys'03)*, pp. 171-180. Los Angeles, November, 2003.

[23] Jeremy Elson. Time synchronization in wireless sensor networks, *Ph.D. dissertation*, May 2003.

[24] S. Ganeriwal, R. Kumar, and M. B. Srivastava, Time-sync protocol for sensor networks, *In Proc. of First ACM Conference on Embedded Networked Sensor Systems (SenSys 2003),* pp. 138-149, Los Angeles, November 2003

[25] Lewis Girod and Deborah Estrin, Robust range estimation using acoustic and multimodal sensing, *In Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001)*, Maui, Hawaii, October 2001.

[26] Nirupama Bulusu, John Heidemann, and Deborah Estrin, GPS-less low cost outdoor localization for very small devices, *IEEE Personal Communications Magazine*, 7(5) pp.28-34, Oct 2000.

[27] Y. Yao and J. E. Gehrke, Query processing in sensors networks, *In Proc. of 1st Biennial Conf. Innovative Data Systems Research (CIDR 2003)*, Asilomar, CA, 2003.

[28] Sanjay Shakkottai , Asymptotics of Query Strategies over a Sensor Network, *In Proc. of IEEE INFOCOM 2004*, Hong Kong, March 2004.

[29] C. Intanagonwiwat, R. Govindan, and D. Estrin, Directed diffusion: A scalable and robust communication paradigm for sensor networks, *In Proc. ACM/IEEE Mobicom 2000*, pp. 56–67, Boston, MA, Aug. 2000.

[30] C. E. Perkins, and Elizabeth Royer, Ad-hoc On Demand Distance Vector Routing, *2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, New Orleans, Louisiana, February 1999.

[31] D. B. Johnson and D. A. Maltz, Dynamic Source Routing in Ad Hoc Wireless Network, *Mobile Computing*, Chapter 5, pp. 153-181.

[32] V. Park and M. S. Corson, A Highly Adaptive Distributed Routing Algorithm for Mobile Wireless Networks, *In Proc. IEEE INFOCOM'97*, Kobe, Japan, April 1997.

[33] J-H. Mo, R.J. La, V. Anantharam, and J. Walrand, Analysis and Comparison of TCP Reno and Vegas, *In Proc. of INFOCOM'99*, New York, NY, March 1999.

[34] K. Fall and S. Floyd, Simulation-based Comparisons of Tahoe, Reno, and SACK TCP, *Computer Communication Review*, V.26 N.3, pp. 5-21, July 1996.

[35] G. Holland and N. Vaidya, Analysis of TCP Performance over Mobile Ad Hoc Networks, *In Proc. of Mobicomm'99*, Seatle, WA, August 1999.

[36] C. E. Perkins and P. Bhagwat, Highly Dynamic Destination Sequenced Distance Vector Routing (DSDV) for Mobile Computers, *In Proc. of ACM SIGCOMM'94*, pp. 234 ~244, Sept. 1994.

[37] T.H. Clausen, G. Hansen, L. Christensen and G. Behrmann, The Optimized Link State Routing Protocol, Evaluation through Experiments and Simulation, *IEEE Symposium on Wireless Personal Mobile Communications*, Sept. 2001.

[38] TinyOS Package, available from http://www.tinyos.net

[39] Mica2 datasheet, http://www.xbow.com/Products/Product_pdf_files/Wireless_pdf/MICA2_Datasheet.pdf

[40] NS-2 simulator, http://www.isi.edu/nsnam

[41] Z. Haas and M. Pearlman, The Zone Routing Protocol (ZRP) for Ad Hoc Networks, *draft-ietf-manet-zone-zrp-00.txt*, IETF MANET Working Group Document, 1997

[42] J. Macker, and M. S. Corson, Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations, *draft-ietf-manet-issues-01.txt*, IETF MANET Working Group, April 1998.

[43] D. D. Clark and D.L. Tennenhouse, Architectural Consideration for a New Generation of Protocols, *In Proc. ACM SIGCOMM'90,* August 1990.

[44] M. Gerla and J.T-C Tsai, Multicluster, Mobile, Multimedia Radio Network, *Wireless Networks 1(3)*, 1995

[45] P. Johansson, T. Larsson, N. Hedman and B. Mielczarek, Routing protocols for mobile ad-hoc network – a comparative performance analysis, *In Proc. of 5th International Conference on Mobile Computing and Networking (ACM Mobicom'99)*, Seattle, Washington, August 1999.

[46] Angin, O., Campbell, A.T., Kounavis, M.E., and Liao, R.R.-F., The Mobiware Toolkit: Programmable Support for Adaptive Mobile Networking, *IEEE Personal Communications Magazine, Special Issue on Adaptive Mobile Systems*, August 1998.

[47] P. Sinha, R. Sivakumar, and V. Bharghavan, CEDAR: a Core-Extraction Distributed Ad hoc Routing algorithm, *In Proc. of IEEE Infocom '99*, New York, NY, March, 1999

[48] S. Ramanathan and M. Steenstrup, Hierarchically-organized, multihop mobile networks for multimedia support, A*CM/Baltzer Mobile Networks and Applications*, Vol. 3, No. 1, pp 101-119.

[49] C. R. Lin and M. Gerla, A Distributed Architecture for Multimedia in a Multihop Dynamic Packet Radio Network, *In Proc. of IEEE Globecom'95*, pp. 1468-1472, Nov. 1995.

[50] IEEE 802.11 Working Group, http://grouper.ieee.org/groups/802/11/

[51] V. Park and M. S. Corson, A Performance Comparison of the Temporally-Ordered-Routing Algorithm and Ideal Link-State Routing, *In Proc. of IEEE Symposium on Computers and Communication '98*, Athens, Greece, June 1998.

[52] S. Das, C. E. Perkins, and E Royer, Performance Comparison of Two On Demand Routing Protocols for Ad Hoc Networks, *In Proc. of IEEE INFOCOM*, Tel Aviv, Israel, March 2000.

[53] The IEEE 802.15 Working Group for WPAN, http://www.ieee802.org/15/

[54] The IEEE 802.16 Working Group on Broadband Wireless Access Standards, http://www.ieee802.org/16/

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

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

[57] R. R-F. Liao and A.T. Campbell, On Programmable Universal Mobile Channels in a Cellular Internet, *In Proc. of 4th ACM/IEEE International Conference on Mobile Computing and Networking (MOBICOM'98)*, Dallas, October, 1998.

[58] C.R. Lin and Mario Gerla, Asynchronous Multimedia Multi-hop Wireless Networks, *In Proc. of IEEE INFOCOM'97*, Kobe, Japan, April 1997.

[59] Global Mobile Information Systems Program, http://www.darpa.mil/ito/research/glomo/index.html.

[60] Linux, RedHat 7.3 distribution, http://www.redhat.com

[61] J. Broch, D. A. Maltz, D. B. Johnson, Y-C Hu, and J. Jetcheva, A Performance Comparison of Multi-Hop Wireless Ad Hoc Network Routing Protocols*,* In Proc. the 4th Annual ACM/IEEE International Conference on Mobile Computing and Networking, Dallas, TX, October 1998.

[62] S. Lu, V. Bharghavan, and R. Srikant, Fair scheduling in wireless packet networks, *In Proc. of ACM SIGCOMM'97*, San Francisco, CA, 1997.

[63] AODV and MACKILL source code distribution, http://user.it.uu.se/~henrikl/aodv

[64] M. Shreedhar and G. Varghese, Efficient Fair Queuing using Deficit Round Robin, *In Proc. of ACM SIGCOMM'95*, Berkeley, California, 1995.

[65] Imad Aad and Claude Castelluccia, Differentiation mechanisms for IEEE 802.11, *In Proc. of IEEE INFOCOM'01* , Anchorage, Alaska, April 2001.

[66] Balachandran A., Campbell A.T., and Kounavis M.E., Active Filters: Delivering Scaled Media to Mobile Devices, *In Proc. of NOSSDAV'97*, St. Louis, MO, 1997.

[67] TORA OPNET source code supplied by V. Park and M. S. Corson, 1997

[68] S-B. Lee, G-S. Ahn, X. Zhang, and A.T. Campbell, INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks, *Journal of Parallel and Distributed Computing, Special issue on Wireless and Mobile Computing and Communications*, Vol. 60 No. 4 pg 374-406, April 2000.

[69] Gomez, J., Campbell, A.T. and H. Morikawa The Havana Framework: Supporting Application and Channel Dependent QOS in Wireless Networks, *In Proc. 7th International Conference on Network Protocols (ICNP'99)*, Toronto, Oct-Nov. 1999

[70] Secure protocols for adaptive, robust, reliable, and opportunistic WINGs (SPARROW) project, available at http://www.cse.ucsc.edu/research/ccrg/projects/sparrow.html

[71] The INSIGNIA Project, http://comet.columbia.edu/insignia/

[72] Ephremides, A. and T. Truong, Scheduling Algorithm for Multi-hop Radio Networks, *IEEE Transaction On Computers*, 38:1353, 1989.

[73] S.B. Lee, G.S. Ahn, and A.T. Campbell, Improving UDP and TCP Performance in Mobile Ad Hoc Networks with INSIGNIA, *IEEE Communication Magazine*, June 2001.

[74] M. Barry, A.T. Campbell, and A. Veres, Distributed Control Algorithms for Service Differentiation in Wireless Packet Networks, *In Proc. of IEEE INFOCOM'2001*, Anchorage, Alaska, April 2001.

[75] Chunhung Richard Lin, On-demand QoS Routing in Multihop Mobile Networks, *In Proc. of IEEE Infocom 2001*, Anchorage, April 22-26, 2001

[76] S. Chen and K. Nahrstedt, Distributed Quality of Service Routing in Ad Hoc Networks, *IEEE Journal on Selected Areas in Communications*, vol. 17, No. 8, Aug 1999.

[77] W.H. Liao, Y.C. Tseng, S.L. Wang, and J.P. Sheu, A Multi-path QoS Routing Protocol in a Wireless Mobile Ad Hoc Network, *Telecommunication Systems* Vol. 19, No. 3-4, pp. 329-347, 2002

[78] G.S. Ahn, A.T. Campbell, A. Veres and L-H Sun, SWAN: Service Differentiation in Stateless Wireless Ad Hoc Networks, *In Proc. IEEE Infocom 2002*, New York, New York, June 2002

[79] T.W. Chen, M. Gerla and J.T. Tsai, QOS Routing Performance in Multi-hop, Wireless Networks, *IEEE International Conference on Universal Personal Communications'97*, San Diego, CA, October 1997.

[80] H. Luo and S. Lu, A Topology-Independent Fair Queuing Model in Ad Hoc Wireless Networks, *IEEE International Conference on Network Protocols 2000*, Osaka, Japan, November 2000.

[81] S. Guo and O.W. Yang, Performance of Backup Source Routing in mobile ad hoc networks, *In Proc. of IEEE Wireless Communication and Networking Conference,* April 2002

[82] A. Nasipuri and S. Das, On-demand multipath routing for mobile ad hoc networks, *In Proc. IEEE ICCCN '99*, Boston, MA, Oct. 1999

[83] M. Garey and D. Johnson, Computer and Intractability: A Guide to Theory of NP-Completeness: W.H. Freeman, 1979

[84] Christine E. Price, Krishna M. Sivalingam, Prathima Agarwal and Jyh-Cheng Chen, A Survey of Energy Efficient Network Protocols for Wireless and Mobile Networks, *In ACM/Baltzer Journal on Wireless Networks*, vol. 7, No. 4, pp. 343 -358, 2001

[85] M. Zuniga, B. Krishnamachari, Analyzing the Transitional Region in Low Power Wireless Links, *In Proc. of IEEE International Conference on Sensor and Ad hoc Communications and Networks*, Santa Clara, CA, Oct. 2004

[86] G. Zhou, T. He, S. Krishnamurthy, and J. Stankovic, Impact of Radio Irregularity on Wireless Sensor Networks, *In Proc. of MobiSys*, MA, June 2004

[87] A. Woo and D. Culler, Transmission Control Scheme of Media Access in Sensor Networks, *In Proc. of ACM Mobicom*, Italy, 2001

[88] C. Wan, S. Eisenman, and A. Campbell, CODA: Congestion Detection and Avoidance in Sensor Networks, *In Proc. of ACM Sensys'03*, CA, Nov. 2003

[89] D. S. J. De Couto, D. Aguayo, J. Bicket and R. Morris. High-Throughput Path Metric for Multi-Hop Wireless Routing, In Proc. of ACM MobiCom, Sept. 2003.

[90] M. Zorzi and R. R. Rao, Geographic Random Forwarding for Ad Hoc and Sensor Networks: Multihop Performance, IEEE Transaction On Mobile Computing, Vol. 2, pp. 337–348, Oct-Dec 2003.

[91] S.-B Lee, G.-S Ahn, X. Zhang and A.T. Campbell, INSIGNIA, IETF Internet Draft, draft-itef-manet-insignia-01.txt, IETF MANET Working Group Document, November 1999.

[92] J. Rosenberg et al., SIP: Session Initiation Protocol, RFC 3261

[93] M. Heissenbuttel, T. Braun, T. Bernoulli, and M. Walchli, BLR: Beacon-Less Routing Algorithm for Mobile Ad Hoc Networks. Elsevier's Computer Comm. Journal, 27(11):1076-1086, July 2004

[94] Matthias Witt and Volker Turau, BGR: Blind Geographic Routing for Sensor Networks. In Proceedings of the 3rd Workshop on Intelligent Solutions in Embedded Systems (WISES'05). Hamburg, Germany, May 2005

[95] M. Zorzi and R. R. Rao, Geographic Random Forwarding for Ad Hoc and Sensor Networks: Multihop Performance, IEEE Trans. On Mobile Computing, vol. 2, pp. 337-348, Oct-Dec 2003

[96] Eugene Shih et al.., Physical layer driven Protocol and Algorithm Design for Energy-efficient Wireless Sensor Networks. In Proc. of ACM Mobicom, Italy, 2001