Click here to view PowerPoint presentation; Press Esc to exit

Interoperable Security, Routing and Quality of Service for Ad-Hoc Network Mobility[1,2,3]

Luiz A. DaSilva, Scott F. Midkiff, Jahng S. Park, George C. Hadjichristofi, Nathaniel J. Davis
Bradley Department of Electrical and Computer Engineering
Virginia Polytechnic Institute and State University
206 North Washington Street, Suite 400
Alexandria, Virginia 22314
USA

{ldasilva, midkiff, jahng.park, ghadjich, ndavis}@vt.edu

Kaustubh S. Phanse
Department of Computer Science and Electrical Engineering
Luleå University of Technology
SE-971 87 Luleå
SWEDEN

Tao Lin
Department of Electrical and Computer Engineering
McMaster University
Hamilton, Ontario L8S 4K1
CANADA

taolin@ieee.org

ABSTRACT

The integration of various network-level functions, including routing, management, and security, is critical to the efficient operation of a mobile ad hoc network (MANET). This paper focuses on network mobility, implying the movement of entire subnetworks with respect to one another, while individual users initially associated with one such subnetwork may also move to other domains. One example is a battlefield network that includes ships, aircraft, and ground troops. In this "network of networks," subnets (e.g., shipboard networks) may be interconnected via a terrestrial mobile wireless network (e.g., between moving ships). We discuss the design and implementation of a new ad hoc routing protocol, a suite of solutions for policy-based

[3] This work was funded in part by the U.S. Office of Naval Research through the Navy Collaborative Integrated Information Technology Collaborative Initiative (NAVCIITI).

DaSilva, L.A.; Midkiff, S.F.; Park, J.S.; Hadjichristofi, G.C.; Davis, N.J.; Phanse, K.S.; Lin, T. (2006) Interoperable Security, Routing and Quality of Service for Ad-Hoc Network Mobility. In *Military Communications* (pp. 21-1 – 21-14). Meeting Proceedings RTO-MP-IST-054, Paper 21. Neuilly-sur-Seine, France: RTO. Available from: http://www.rto.nato.int/abstracts.asp.

network management, and approaches for key management and deployment of IPsec in a MANET. These solutions are integrated with real-time middleware, a secure radio link, and a topology monitoring tool. We briefly describe each component of the solution and focus on the challenges and approaches to integrating these components into a cohesive system to support network mobility. We evaluate the effectiveness of the system through experiments conducted in a wireless ad hoc testbed.

1.0 INTRODUCTION

There has been significant research on mobile ad hoc networks in recent years. To date, most research has focused on a single aspect of the problem, such as medium access, routing or mobility support. The focus of this paper is on the *integration* of related functions, including network management, quality of service (QoS), routing, and security to support mobile ad hoc networks. In particular, we consider network mobility, rather than node mobility, implying the movement of entire subnetworks with respect to one another, while individual users initially associated with one such subnetwork may also move to other domains. One example is a battlefield network that includes ships, aircraft, and ground troops. In this "network of networks," subnets (e.g., shipboard networks) may be interconnected via a terrestrial mobile wireless network (e.g., between moving ships). Mobile users are initially associated with their home networks, but are free to move between domains. Challenges in such a scenario include interoperation among different platforms, maintenance of security associations, and distribution of policies to preserve quality of service.

Figure 1 summarizes the aspects of network integration considered in our work. We propose a modification of the Open Shortest Path First (OSPF) routing protocol (1) that uses a minimum connected dominating set (MCDS) of nodes to propagate route updates. Security (2) is accomplished through the tunneling of data over the ad hoc network using Internet Protocol Security (IPsec) and Generic Routing Encapsulation (GRE). Authentication keys are dynamically distributed to the network nodes using multiple key repositories. To achieve quality of service (3), bandwidth is allocated according to a distributed policy-based network management mechanism. Some nodes in the network have the capability to perform topology monitoring (4) through periodic exchange of Simple Network Management Protocol (SNMP) messages. To support real-time applications, some hosts run middleware (5) responsible for identifying deadline requirements of the application associated with utility functions and marking packets accordingly using the DiffServ Code Point (DSCP) field of the IP header. Finally, a secure radio link (6) is provided for some of the links in the network.

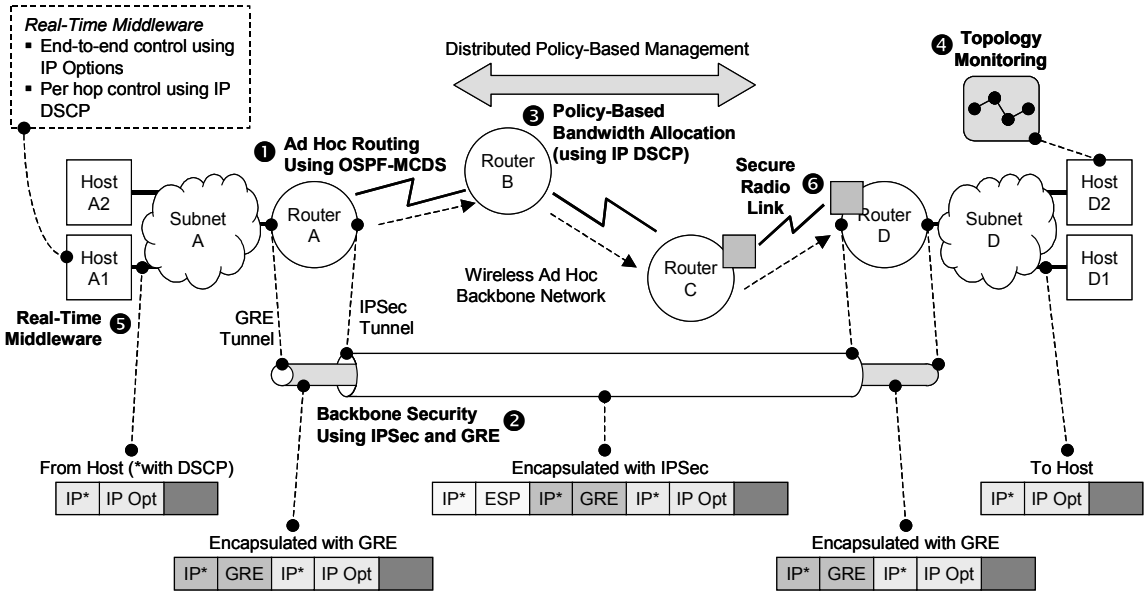

Figure 1: Integration of network management, routing, QoS and security in a MANET.

In this paper, we present novel algorithms and protocol extensions for routing (Section 2.0), network management (Section 3.0), and security (Section 4.0) in MANETs. All of these protocols have been prototyped and tested in a wireless network testbed, as described in Section 5.0. We also describe challenges and solutions in integrating these mechanisms to form a cohesive suite of solutions in support of preserving reliability and quality of service in ad hoc networks. We conclude by discussing major lessons learned and directions for future research in Section 6.0.

2.0 ROUTING

A number of routing protocols have been proposed for MANETs, including Ad-hoc On-demand Distance Vector (AODV) [1], Dynamic Source Routing (DSR), Optimized Link State Routing (OLSR) [2], and Topology Broadcast Based on Reverse Path Forwarding (TBRPF) [3]. AODV and DSR, both reactive routing protocols, cannot always provide shortest-path routing since they do not update a route in use unless the route is broken due to the mobility of network components. Reactive protocols may also present high control overhead when a large number of traffic flows are present [4]. Besides these potential disadvantages, reactive protocols do not provide full topology information, which might be required by a network management application, such as the policy-based management system described in the next section. Proactive routing protocols, including OLSR and TBRPF, do provide shortest-path routing and more extensive topology information, but at the cost of high control overhead for topology advertisements. In particular, TBRPF allows the broadcast of full topology information, but may produce redundant control traffic since a node may receive the same link state information from multiple neighbors. We propose and implement a proactive routing protocol that locally maintains full topology information and, also, imposes low control overhead [5].

Our proposed protocol, OSPF-MCDS [5], is similar to OSPF, a widely used routing protocol designed for wired networks. We replace the concept of designated routers in OSPF by a minimal connected dominating set of routers and simplify control messages. A connected dominating set (CDS) is a set of routers that forms a connected topology with the property that any other router that is not in the set has at least one neighbor in the

set. Figure 2 illustrates how OSPF-MCDS works. The set of black nodes in the figure is chosen as an MCDS. Only nodes in this set will forward any broadcast topology control messages. For example, when the link between nodes 1 and 4 becomes available, one of the end nodes, say node 1, first broadcasts the existence of this new link. The link state information is then propagated to other nodes via nodes 3, 5, and 6. By the definition of a CDS, broadcast topology control messages can reach all nodes in the network. Thus, all nodes maintain identical copies of the network topology (except for short-term inconsistencies due to delays in the propagation of control messages) and build their own shortest path trees and generate routing entries accordingly. Unlike some other protocols that use CDS nodes as default gateways for routing, e.g. OLSR [2], the Core Extraction Distributed Ad-hoc Routing (CEDAR) protocol [6], and the simple gateway protocol proposed by Wu and Li [7], OSPF-MCDS can generate smaller CDSs and only uses CDS nodes to broadcast topology information. Relay nodes in OSPF-MCDS are selected only to propagate control messages. They do not necessarily serve as gateway routers for user data packets, unlike in OLSR, where relay nodes are chosen as gateways for user data packets. When the traffic load is heavy, using CDS nodes as gateways may increase collisions between data packets and control packets, a potential problem in OLSR, CEDAR, and Wu and Li's simple gateway protocol.

Broadcast using an MCDS can reduce the number of retransmissions compared with blind broadcast (where all nodes rebroadcast the control messages that have not been received before) and, thus, achieves the goal of low control overhead. The redundant traffic eliminated by using a CDS is proportional to the number of non-CDS nodes divided by the total number of nodes in the network. A simple simulation is presented here to illustrate the improvement [5]. In the simulation, n nodes are randomly placed in a 100×100 square unit area. Radio range determines connectivity between two nodes. If radios are capable of longer transmission and reception ranges (for instance, by increasing power or antenna gain), more links are viable, resulting in a more densely connected network. Three radio ranges, 25, 50, and 75 units, are used. All possible node sets are examined to find an optimum CDS for all topologies. The CDS with the minimum size is kept. For each set of parameters, we replicate the experiment 1000 times with different random node placements. The graph in Figure 2 shows the percentage of overhead reduced using a CDS compared with blind broadcast. Overhead is reduced by over 50% for all radio ranges and values of n. Savings increase when radio range increases, implying greater benefit in dense networks. Besides the advantage of low control overhead, OSPF-MCDS also maintains shortest-path routes and can provide full topology information. The link costs can optionally be defined according to traffic load or power consumption for load balancing or power efficient routing.

Using the MCDS concept to reduce control overhead is a subject of current research. The algorithm we use in OSPF-MCDS exhibits better performance compared to other known approaches in terms of the average size of CDSs, which, in turn, determines the number of retransmissions of control messages and the control overhead [4][5]. A simulation study [4] also demonstrates that OSPF-MCDS has low overhead compared to reactive protocols, such as AODV, especially when the number of traffic flows is large.

In our integrated testbed, described in Section 5.0, a copy of OSPF-MCDS runs in every gateway node. It maintains a local routing table to enable subnet-to-subnet routing. Moreover, it provides hop counts between any pair of nodes to the policy-based management system, which is discussed in the next section.

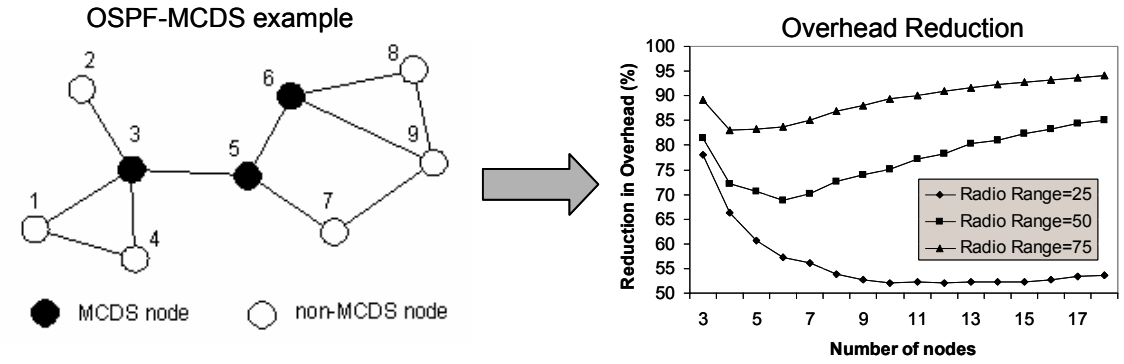

Figure 2: An example MANET running OSPF-MCDS.

3.0 POLICY-BASED QUALITY OF SERVICE

Unlike legacy network management, which generally involves configuring and managing each network entity individually, Policy-Based Network Management (PBNM) configures and controls the network as a whole, providing the network operator with simplified, logically centralized and automated control over the entire network. PBNM can be used to control different networking capabilities such as quality of service, network security, access control, and dynamic IP address management. PBNM provides a viable solution for managing mobile ad hoc internetworks – a consortium of multiple subnetworks controlled by distinct organizational policies. We propose a solution suite [8] to apply the policy-based approach to managing QoS in MANETs. The four components of this suite are briefly described below.

- *k-hop cluster management:* Using clustering, we limit the number of hops between a policy server and its clients. We propose two ways to implement clustering: (i) by taking advantage of the topology information gathered by the underlying proactive ad hoc routing protocol, whenever such information exists; and (ii) through interaction between the Common Open Policy Service (COPS) protocol based application layer and the IP layer, the idea being to control the time-to-live (TTL) field in the IP header for the COPS Keep-Alive (KA) messages exchanged periodically by the policy server and client. Both methods enable clustering with minimal additional overhead.
- *Dynamic Service Redundancy (DynaSeR):* The DynaSeR solution implements redirection and delegation that allow the PBNM system to improve its service coverage. Redirection is a server-centric way of helping a client leaving its current cluster to discover a new server, while delegation allows dynamic invocation of policy server instances on demand to cover as many clients in the network as possible by covering those that lie outside all the existing clusters. We extend the standard COPS for PRovisioning (COPS-PR) protocol, adding delegation capabilities.
- *Service discovery:* We implement a lightweight service discovery mechanism to facilitate automated discovery of policy servers in the network. Two types of messages are used: Service Advertisement (SA) and Client Service Request (CSRQ). A policy server periodically advertises itself via a limited *k*-hop broadcast of SA messages. A client that does not receive an SA message within a certain time interval broadcasts a CSRQ message. The server that may have moved within *k* hops of the client responds with a unicast SA message. Alternatively, a client node that is currently being serviced, upon hearing a CSRQ message, may volunteer to act as a delegated server.

- *Inter-domain policy negotiation:* We extend the COPS-PR protocol to facilitate inter-policy server communication and to support policy negotiation between different domains. This allows seamless QoS provisioning for nodes moving across different domains in a mobile ad hoc internetwork.

We implement our proposed schemes and protocols both as a prototype in a Linux-based ad hoc network testbed, as discussed later, and as simulation models in QualNet. The PBNM system prototype is integrated with the OSPF-MCDS proactive ad hoc routing daemon to implement *k*-hop clustering and its operation is demonstrated over a heterogeneous (wired and wireless) ad hoc network secured using IPsec and GRE tunneling. The effectiveness of the PBNM system for managing QoS is illustrated using soft real-time applications [9]. Almost seamless QoS is obtained for real-time applications hosted on a mobile device moving across an emulated multi-domain ad hoc network.

Through simulation, we study the service availability and overhead of the PBNM system as a function of mobility, network density, and cluster size. We adopt the random waypoint mobility model to simulate node mobility. Our proposed management solution is found to scale well (up to 100 nodes were considered). The tradeoff lies in increased predictability and reliability for small cluster sizes versus improved service availability for large cluster sizes. Our proposed delegation scheme addresses this trade-off and allows the PBNM system to improve its service coverage while maintaining smaller cluster sizes. As shown in Figure 3, delegation improves the policy service availability by up to 25%. Thus, we can generally use small clusters for localized management, while catering on demand to client nodes that fall outside existing clusters. A complete set of results is provided in [10].

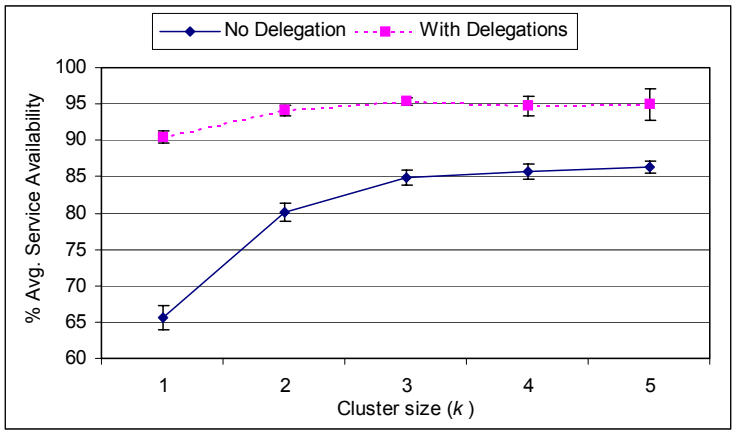

Figure 3: Improvement in service availability through the use of delegation.

4.0 SECURITY

In the area of security, we focus on the interoperability of IPsec and key management over multiple platforms, including Cisco, Microsoft Windows 2000, and Red Hat Linux, with different emerging technologies such as OSPF-MCDS, QoS, and real-time systems. FreeS/WAN IPsec, a freely available commercial-off-the-shelf implementation of IPsec, is installed in all gateway nodes. The selection of FreeS/WAN is based on the availability of IPsec implementations for RedHat Linux and functionality. FreeS/WAN IPsec was the only version available at the time of the testbed deployment. Even though there is an IPsec implementation built into the latest RedHat Linux kernel, that implementation lacks the functionality of opportunistic encryption that is used in our testbed.

To deploy a security mechanism, such as IPsec, in a network, two peers must have a preconfigured level of trust between them. This level of trust is achieved via authentication. These keys or certificates can be distributed to the nodes automatically via a key management system. Key management entails the secure generation, distribution, revocation, re-issuance and storage of keys on network nodes. In our work, we address the storage and distribution aspects of key management. We also investigate ways of providing redundancy and robustness for key management to facilitate the establishment of IPsec security associations in a MANET and propose a key management system for such an environment.

Key negotiation in our system is provided using automatic keying via the Internet Key Exchange (IKE) protocol [11]. Authentication is achieved using asymmetric keys, which are easier to handle than symmetric keys since ownership of public keys does not compromise security. The asymmetric keys are installed in multiple key distribution centers. A relatively new feature of IPsec implemented in FreeS/WAN IPsec, known as opportunistic encryption, allows this functionality, which is well suited for the dynamic topology of a MANET. Opportunistic encryption enables any two systems to authenticate each other without requiring a pre-shared key negotiated out of band. The public keys of the nodes are stored on a Domain Name Service (DNS) server, which removes the need to set up the keys in the configuration file and decreases key management overhead. The DNS servers are set up in different subnets, so that they are protected by the IPsec gateways. The DNS servers are implemented using BIND in Linux. Once communication with any peer is established, nodes can dynamically obtain each other's public key during the IKE negotiation and set up security associations between them. A disadvantage of opportunistic encryption is that it is currently vulnerable to a man-in-the-middle (MITM) attack. The use of secure DNS using DNS security extensions (DNSSEC) may address this vulnerability. The interoperation of DNSSEC features with IPsec is an area of future work.

The proposed key management system also implements certificate issuance and maintenance. It differs from existing systems because it dynamically switches from a centralized scheme of trust distribution to a more distributed scheme, which is better suited for MANETs. Authentication is achieved via asymmetric keys embedded in Certificate Authority (CA) certificates. CA certificates offer the advantage of identifying the user as well as the IP address of a node, thus removing the need for dual authentication per host. The nodes are also assigned different levels of trust by the key management system, accounting for the fact that not all nodes in a network have the same trustworthiness.

The key management system uses a modified hierarchical model as shown in Figure 4. The Root Certificate Authority (RCA) is assumed to be off-line. Any node that has an RCA certificate obtained via out-of-band methods can act as a Delegated Certificate Authority (DCA). Thus, the key management system requires minimal pre-configuration of trust for the nodes. The DCAs have the responsibility of issuing, distributing, revoking and storing certificates of nodes. Furthermore, any node in the network that is not a DCA can assume the role of a Temporary Certificate Authority (TCA) and sign temporary certificates for co-located nodes.

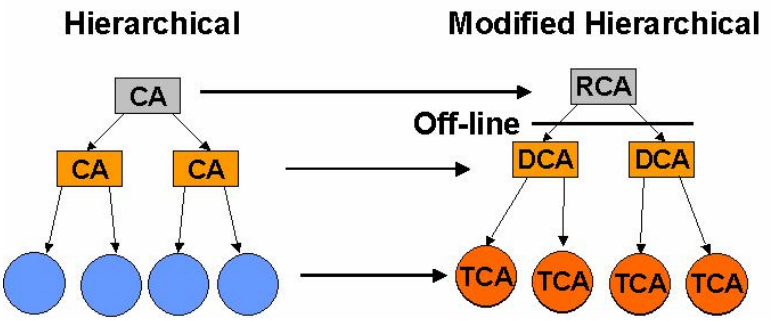

Figure 4: The key management system adopts a modified Public Key Interchange (PKI) model.

Service availability is increased in a number of ways. The system offers multiple DCAs that generate, deposit, reissue, revoke, and distribute certificates to the nodes. If all the DCAs are unavailable, a node can obtain a peer's certificate from any node that already trusts that peer. This functionality is achieved by having each node store the certificates of the nodes that it trusts. Furthermore, the system decreases the frequency of certificate issuance and revocation by relaxing time constraints. Certificates are reissued whenever a node or DCA desires and are revoked whenever a node is compromised.

This system does not necessarily require out-of-band authentication with a DCA. New nodes joining the network can simply register at a lower trust level with the DCA if they are unable to authenticate using out-of-band methods. In this way, they are motivated to register with out-of-band methods as soon as they can communicate with a DCA. In addition, the key management system is flexible enough to accommodate new nodes when the DCA is unavailable. New nodes that join the network and are preconfigured with an RCA-certificate can temporarily establish trust with other nodes. If they do not possess a certificate, they can obtain a temporary certificate from any of the TCAs that are physically co-located by first authenticating out of band. As a result, they can temporarily be accepted into the network until they can register at a DCA. The key management system maintains sufficient levels of security by combining node authentication with an additional element, node behavior. A behavior-grading scheme allows each node to grade the behavior of other nodes. The key management system records and evaluates the behavior of nodes and provides credentials to negotiating peers for deciding whether they should trust each other. The effectiveness of the proposed key management in distributing trust is a subject of ongoing research.

The subnetworks in our "network of networks" communicate with each other via secure tunnels. The different configurations that can be used to achieve this functionality are either tunnel mode IPsec or transport mode IPsec with GRE tunnels. Transport mode IPsec with GRE tunnels is not used because IPsec does not properly configure routing for the IPsec virtual interfaces when path lengths between nodes in the same subnet are greater than one. As a result, packets from one node cannot be sent to another node via peer nodes, unless those two nodes are directly connected. Therefore, tunnel mode IPsec is used instead of transport mode with GRE tunnels.

The real-time middleware with which we operate [9] makes use of the IP options field in the IP header to encode deadline information and current latency experienced by the datagram. However, the FreeS/WAN IPsec implementation drops packets that utilize IP options in tunnel mode, not complying with RFC 2401 [12]. To preserve the IP options field and allow the real-time system to interoperate with IPsec, GRE is used in conjunction with IPsec. GRE tunnels encapsulate any network layer protocol unit, allowing its transmission over any other network layer protocol. To use GRE with IPsec, GRE tunnels are attached to the private side of the gateways so that the source and destination addresses of the packet comply with the IPsec policy.

Interoperability of IPsec with QoS schemes is also achieved by setting both the IPsec and GRE protocols to preserve the DSCP field in the IP header through the different levels of encapsulation. The overhead impact of GRE is an additional 24 bytes per IP packet.

Special steps must be taken to integrate MANET routing protocols with IPsec. FreeS/WAN IPsec creates a virtual interface for an IKE negotiated tunnel so that packets can be routed through that interface. One of the limitations of this implementation is that it uses routing to determine the IPsec policy that should be applied to every packet. More specifically, packets destined for a particular subnet and requiring encryption have to be routed through the corresponding IPsec virtual interface for IPsec to be applied to those packets. Furthermore, MANET routing protocols modify the subnet routing entries based on dynamic topology changes. These modifications introduce interoperability issues because the IPsec virtual interface and the corresponding subnet routing entry have the same network mask. A solution to this conflict that allows IPsec to be deployed in a MANET is to assign a higher subnet mask to the IPsec interface. Thus, the subnet traffic is directed through the IPsec interface complying with the IPsec policy and MANET routing does not interfere with the IPsec virtual interface. This method decreases the size of the subnet behind the gateway and increases the number of possible subnets. A more complete and robust solution for IPsec interoperation with MANET routing requires modifications to the IPsec implementation so that IPsec is independent of routing in the Linux kernel. We have not implemented this more general solution.

In addition to security provided by IPsec, we incorporate secure radio links developed by Virginia Tech's Configurable Computing Laboratory [13]. The secure radio links are secure configurable platforms that resist reverse engineering, thus protecting both the data and the intellectual property contained in them.

5.0 INTEGRATION AND TESTBED

In this section, we describe the integration of the mechanisms described above into the wireless ad hoc network testbed illustrated in Figure 5. Gateways G1 through G7 are interconnected via a "dynamic switch." The dynamic switch emulates a mobile wireless topology, including packet loss and constrained capacity [14]. It allows repeatable, controlled experiments in a MANET environment with many nodes in a limited test bed area. The figure shows a particular wireless topology. By changing the switching table of the dynamic switch, Gateways G1 through G7 can be logically reconnected to form different topologies. The operation of the dynamic switch is transparent to each node. The nodes are stationary and connected by wires, but the protocols and applications running on the nodes behave as if they were in a MANET environment.

Whatever the topology may be, the connectivity of the network is maintained by the OSPF-MCDS routing protocol discussed in Section 2.0. OSPF-MCDS runs on each gateway, maintaining connectivity and ensuring the correct routing of packets with low overhead. A topology monitoring tool developed as part of this effort provides a real-time graphical view of the topology and the connectivity of the gateways. A connection between any pair of gateways can be secured by using IPsec/GRE tunnels as discussed in Section 4.0. The servers and clients of the policy-based network management scheme, described in Section 3.0, take advantage of the efficient routing protocol and the secure connectivity to provide differentiated services, in terms of allocated bandwidth, to different applications.

Next, we describe three test scenarios to examine the correct operation of the different protocols and the integration of these protocols. Scenario 1, shown in Figure 6, tests the performance of the OSPF-MCDS routing protocol and the PBNM scheme. It uses true wireless mobile nodes. Gateway G12 is initially connected to Gateway G9 with bandwidth reservation that ensures a high level of QoS. As Gateway G12 moves towards Gateway G10 (and away from Gateway G9), OSPF-MCDS detects a new link between

Gateways G10 and G12, updates the topology, and maintains the connectivity. At the same time, the policy server at Gateway G10 communicates with Gateway G9 to provide the same level of QoS that Gateway G12 was receiving from Gateway G9. To visualize the effects of link loss, reestablishment of the link, and QoS allocation, we transmit a video image from Gateway G12 to Gateway G6 via Gateway G9 initially and then via Gateway G10. The quality of received video stream via Gateway G10 is initially poor, but as soon as the policy is negotiated, the video stream quality improves, as illustrated in the graphic in Figure 6.

Scenario 2, also depicted in Figure 6, tests the network security capabilities of the test bed.A host connected to Gateway G9 receives HyperText Transfer Protocol (HTTP) packets from an HTTP server in the subnet behind Gateway G1. Without the IPsec tunnel between Gateway G1 and Gateway G9 (via Gateway G6), a hostile packet sniffer can capture and decipher data packets over the wireless link between Gateway G6 and Gateway G9. An IPsec tunnel between Gateway G1 and Gateway G9 is established using IKE. During the IKE negotiation the authentication keys are dynamically obtained from any of the available DNS servers (hosts S1 or S3 in Figure 5). Once the nodes are authenticated and IPsec is deployed, the hostile sniffer can no longer decipher the captured packets.

Scenario 3 tests the integration with real-time middleware. Application packets are transmitted from subnet hosts of Gateway G1 and Gateway G9 (hosts S1 and S9a) to subnet host S2 of Gateway G2, as shown in Figure 5. These packets are beneficial to host S2 only if they arrive within the deadlines indicated by the time-utility functions marked on each packet. The policy server at Gateway G7 and clients at Gateways G4 and G6 limit the bandwidth used by background traffic and allocate sufficient bandwidth so that the application packets do not miss their deadlines. The topology and routing are provided by the OSPF-MCDS routing protocol and the channels between Gateways G1 and G2 and Gateways G9 and G2 are secured by IPsec tunnels. Almost seamless QoS is observed for real-time applications transmitted from hosts S1 and S9a to host S2.

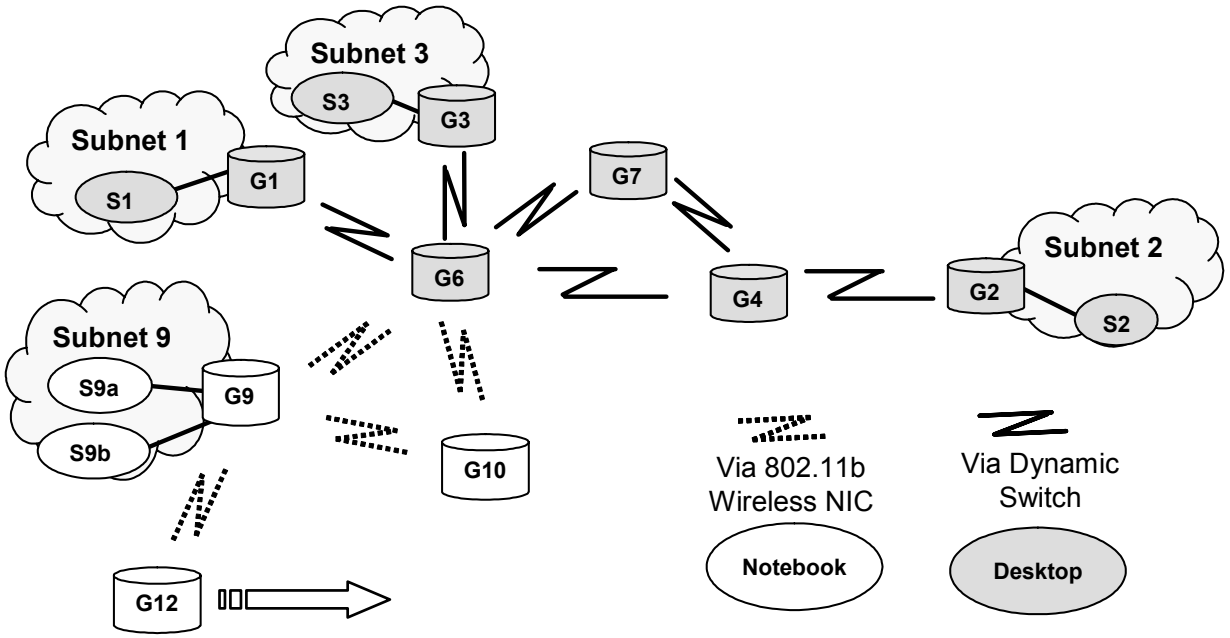

S: Subnets
G: Gateways

S1: RT traffic source 1
S2: RT traffic destination
S9a: RT traffic source 2
S9b: HTTP client (Windows machine)

G1: HTTP server
G4: Policy client (Demo 3)
G6: Vic receiver; Policy client (Demo 3)
G7: Policy server (Demo 3)
G9: Vic router; Policy server (Demo 1)
G10: Vic router; Policy server (Demo 1)
G12: Vic source (with camera)

Notes: G6 has one wired interface and one wireless interface. G9/S9a/S9b, G10, and G12 are placed on carts for mobility experiments.

Figure 5: Wireless network testbed.

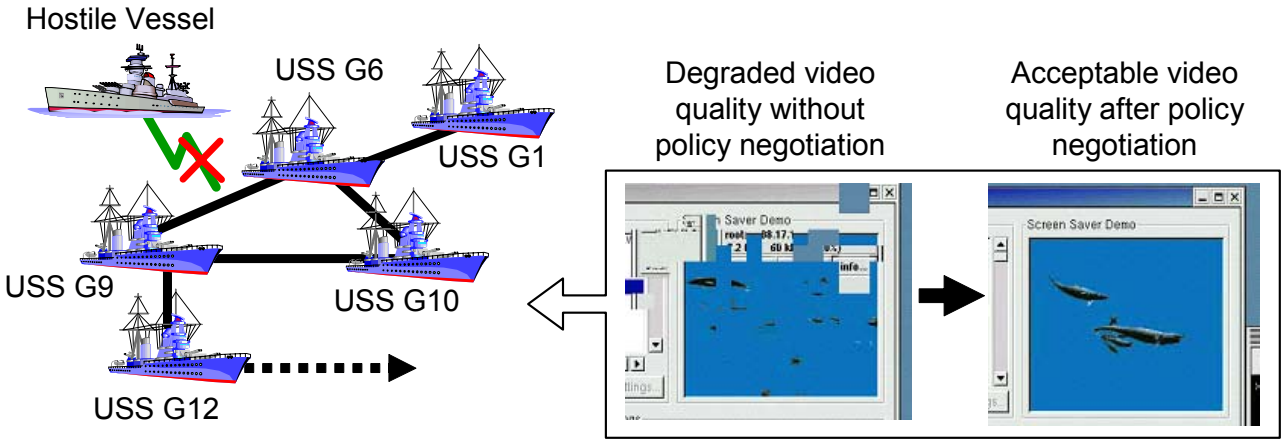

Figure 6: OSPF-MCDS, PBNM, and network security test scenarios.

6.0 CONCLUSIONS

As mobile ad hoc networks mature, it is necessary to integrate the various mechanisms and protocols that have been advanced into a cohesive system that supports reliable, secure communications and quality of service in this dynamic environment. In this paper, we presented solutions for:

- Routing in the mobile backbone using OSPF-MCDS, based on the widely used OSPF routing algorithm, to support wireless interfaces and improve performance in a wireless mobile environment;
- Management of bandwidth allocation using a decentralized PBNM scheme;
- Secure tunnels between subnet gateways using IPsec and GRE in a manner that is integrated with the routing and policy-based network management schemes; and
- Integration of PBNM with real-time middleware by using scheduling at hosts within a subnet running the real-time middleware and supporting modified IP Differentiated Services (DiffServ) in the backbone network.

The integration of the various functions we describe here was not without its challenges, especially since most of the software consisted of working prototypes. Significant work went into fixing bugs as the integration proceeded. Another difficulty was the unreliable or unexpected behavior of IEEE 802.11b connections when we tested the routing protocol. The signals were sensitive to the number people between nodes and their movement, making it difficult to obtain consistent data in different repetitions of each experiment. This experience emphasized the importance of a topology emulator like the dynamic switch described here for wireless test beds. Without it, the integration would have taken much longer (and caused much more frustration).

Support for real-time applications requires tight integration between the policy-based QoS management, security and routing functions. For instance, the policy server's need to obtain topology information had to be considered during implementation of the OSPF-MCDS prototype. Further, we use GRE tunnels to facilitate the transport of real-time traffic, whose QoS requirements are indicated using the IP options field, in IPsec tunnels. Proper configuration of the IPsec and GRE tunnels is required to ensure that the DSCP field is copied from the inner IP header to the outer IP header.

Lessons learned while investigating the security aspects using the testbed helped us to assess the maturity of relevant technologies. Even though IPsec is superior for this application compared to other security systems, such as SSL, it offered limited functionality and flexibility. The integration of IPsec with the various other system components required a number of adjustments to obtain the desired functionality. Some of the difficulties were due to deviation of the FreeS/WAN implementation from the IPsec architecture, as stated in RFC 2401, in conjunction with FreeS/WAN implementation limitations. Additional difficulties were due to the inability to utilize security policies and assess the state of the security associations and the need to use dual-authentication in multi-user gateways. Different mechanisms proposed in Internet drafts will likely increase the acceptance of IPsec. These include an IPsec flow monitoring Management Information Base (MIB), an IPsec Policy Information Base (PIB) [15], and an IPsec information policy configuration model. However, fully functional implementations will likely not be available in the immediate future.

Current work being undertaken as part of this project includes an experimental study of inter-operation among different MANET routing protocols, an investigation of the proposed key management system with respect to both functionality and security, analytical modeling of the proposed PBNM system using stochastic Petri nets, and an extension of the management system for distributed key management.

7.0 REFERENCES

[1] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc On Demand Distance Vector (AODV) Routing," IETF RFC 3561, July 2003.

[2] T. Clausen and P. Jacquet, eds., "Optimized Link State Routing Protocol (OLSR)," IETF RFC 3626, October 2003.

[3] R. Ogier, F. Templin, and M. Lewis, "Topology Dissemination Based on Reverse-Path Forwarding (TBRPF)," IETF RFC 3684, February 2004.

[4] T. Lin, S. F. Midkiff, and J. S. Park, "A Framework for Wireless Ad Hoc Routing Protocols," *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, vol. 2, March 2003, pp. 1162-1167.

[5] T. Lin, S. F. Midkiff, and J. S. Park, "Approximation Algorithms for Minimal Connected Dominating Sets and Application for a MANET Routing Protocol," *Proc. IEEE International Performance Computing and Communications Conference (IPCCC)*, April 2003, pp. 157-164.

[6] P. Sinha, R. Sivakumar, and V. Bharghavan, "CEDAR: Core Extraction Distributed Ad Hoc Routing," *Proc. IEEE INFOCOM*, March 1999, pp. 202-209.

[7] J. Wu and H. Li, "A Dominating-Set-Based Routing Scheme in Ad Hoc Wireless Networks," *Telecommunication Systems Journal*, vol. 18, no. 1-3, September-November 2001, pp. 13-36.

[8] K. Phanse, "Policy-Based Quality of Service Management in Wireless Ad Hoc Networks," Ph.D. dissertation, Virginia Polytechnic Institute and State University, August 2003.

[9] K. Channakeshava, "Utility Accrual Real-time Channel Establishment in Multi-hop Networks," M.S. thesis, Virginia Polytechnic Institute and State University, August 2003.

[10] K. Phanse and L. A. DaSilva, "Protocol Support for Policy-Based Management of Mobile Ad Hoc Networks," *Proc. IEEE/IFIP Network Operations and Management Symposium (NOMS)*, April 2004, pp. 3-16.

[11] P. Hoffman, "Internet Key Exchange (IKE) Monitoring MIB," IETF, draft-ietf-IPsec-ike-monitor-mib-04.txt, April 2003.

[12] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol," IETF RFC 2401, November 1998.

[13] R. J. Fong, S. J. Harper, and P. M. Athanas, "A Versatile Framework for FPGA Field Updates: An Application of Partial Self-reconfiguration," *Proc. 14th IEEE International Workshop on Rapid Systems Prototyping*, June 2003, pp. 117-123.

[14] T. Lin, S. F. Midkiff, and J. S. Park, "A Dynamic Topology Switch for the Emulation of Wireless Mobile Ad Hoc Networks," *Proc. IEEE Conference on Local Computer Networks (Workshop on Wireless Local Networks)*, November 2002, pp. 791-798.

[15] M. Li, D. Arneson, A. Doria, J. Jason, C. Wang, and M. Stenberg, "IPsec Policy Information Base," IETF, draft-ietf-ipsp-IPsecpib-08.txt, May 2003.