Using Modelling and Simulation to Evaluate Network Services in Maritime Networks

David Kidston
*Communications Research Centre
Canada*

1. Introduction

Maritime networks are composed of a number of mobile and static nodes that have some intermittent wireless connections, as is typical of mobile ad-hoc networks (MANETs), but also have steady (satellite) links and are continuously powered, as is typical of fixed wired networks. Combined, these characteristics provide unique operational challenges not conducive to the use of existing MANET or fixed network techniques. Since maritime units operate in a low bandwidth environment, the efficiency of network services is critical. The lack of power constraints and slower mobility require a less dynamic solution than that required for MANETs. The problem we deal with in this chapter is how to manage the Quality of Service (QoS) achieved by application traffic while optimising the use of semi-reliable and limited-capacity links. This is a traffic engineering (TE) problem.

As part of a research effort to provide enhanced communications capabilities in a maritime network, we proposed and then investigated a number of network services using the OPNET discrete event simulation (DES) tool. The modelling of this type of network provided some unique challenges. The combination of link types has not been previously described in the literature and existing link models had to be customised for their unique low bandwidth characteristics. Similarly, routing in this kind of mobile environment has not been studied. Finally, there has been limited work on modelling the traffic characteristics of maritime networks.

We begin this chapter with a description of the network, mobility, and traffic models developed to simulate the maritime environment. This is followed by a description of four network services that were designed to aid in network management and provide improved QoS in maritime networks. The first service is a traffic monitoring service that matches the amount of traffic it produces with its knowledge of the current load of the network. Second, a traffic prioritisation service uses weighted fair queuing (WFQ) to prioritize traffic in the maritime environment based on dynamically assigned priorities. Third, an adaptive routing service uses multi-path labelled switching (MPLS) to divert traffic from overloaded links. Fourth, we describe our resource reservation service (RRS), a distributed admission control protocol designed to provide some guarantees of end-to-end bandwidth for critical traffic in the maritime environment. The RRS includes a number of features specifically tailored for maritime networks including multi-route probing, aggregated pre-emption, and improved

robustness. In our simulations, these services were found to provide network awareness and significantly improve the timeliness of prioritised flows.

This chapter continues with a description of the results of our simulations, which we gathered based on a process based on five criteria developed to provide improved credibility. The chapter ends with a review of the limited related work in this area followed by a number of conclusions and a discussion of possible future work in this area.

2. Maritime Networks

We based our models of a maritime network on a naval task fleet deployment. In such deployments, a relatively small number of nodes (ships) are dispatched as a group. This is commonly between 2 and 5 nodes (AUSCANZUKUS, 1999). In addition, one or more shore stations provide most server-based application services and act as a satellite switching centre. This environment may also be applicable to a commercial enterprise such as a shipping company or emergency operations at sea such as coast guard duties, though these alternatives have not been investigated. A typical maritime network is shown in Figure 1.

Fig. 1. Typical Maritime network

Maritime networks thus consists of a Network Operation Centre (NOC) which acts as a land based relay for all satellite communication, a limited number of mobile nodes (ships or potentially maritime land/air units), and the bearers that connect them. A commercial

satellite ground station may also be included. The features of maritime networks critical to our network model are expanded upon in the following subsections.

2.1 Communication Bearers

A maritime network is composed of a variety of strategic and tactical communications links. The communications bearers available to transfer information within the network include: commercial satellite (e.g. INMARSAT B), ship-to-shore satellite networks (e.g. SHF SATCOM), High Frequency (HF) extended and beyond line-of-sight (HF ELOS/BLOS) and UHF/VHF line-of-sight (UHF/VHF LOS). A sample of the communications types and capabilities from (AUSCANZUKUS, 2003) are given below in Table 1.

Link Type	Rate	Use
UHF/VHF LOS Radio	Shared 64 kbps	Main data bearer for ship-to-ship communication. Used over short distances (20-50 Nm)
HF BLOS Radio	4.8-9.6 kbps	Email, chat, low data rate apps. HF Sky wave (2000-3000 Nm)
HF ELOS Radio	4.8-9.6 kbps	Email, chat, low data rate apps. HF Surface wave (200-300 Nm)
INMARSAT B Satellite	64 kbps	Main satellite connection for most ships - point to point data bearer
SHF Satellite	Up to 512 kbps	High capacity satellite - point to point data bearer.
25Khz UHF Satellite	Up to 48 kbps	Low bandwidth satellite with limited IP capability (Email, chat, low data rate apps)
5Khz UHF Satellite	Up to 9.6 kbps	Low bandwidth satellite (Email, chat)

Table 1. Communications Subnet Matrix

Mobile maritime nodes (ships) most commonly communicate using a combination of two modes. First, they may communicate back to their strategic network operation centre (NOC) using satellite communications (e.g. INMARSAT, SHF SATCOM). Satellite communications can also be relayed at the NOC to provide indirect ship-to-ship communications. Satellite communications provide high bandwidth but high delay and high cost communications. Second, ships communicate directly with other ships via limited range radio (e.g. UHF/VHF LOS). Recently UHF/VHF relay technology has improved to the point that LOS radio systems may form mobile ad-hoc networks (MANETs) (Jorgensen et al. 2005). These networks provide low cost, low bandwidth and low delay communications over a limited distance.

2.2 Routing Capabilities

Naval maritime networks are now IP-based (AUSCANZUKUS, 2003). By default, the network topology is driven by the routing protocols used to achieve connectivity (e.g. OSPF). Each network is typically divided into separate Autonomous Systems (AS). In

maritime networks an AS is a collection of mobile nodes and shore station nodes connected by a collection of backbone subnets. The shore-stations may be gateways to a third party backbone WAN (e.g. Internet Service Provider) or to another military WAN.

Routing in this environment currently relies on OSPF within an AS with the link cost metric set to increase with decreasing bandwidth (Holliday, 2005). This means that the link with the highest bandwidth is used to the exclusion of any other links that may be available. Due to its high bandwidth, SATCOM will be used predominantly. When low-bandwidth LOS links are the only links available, they are often overloaded with high bandwidth traffic. Between autonomous systems, BGP4 is used, though we are currently assuming a single autonomous system (thus a single OSPF area).

As mentioned in Section 2.1, technology has been developed to allow maritime nodes to form a MANET from their available LOS bearers. Combining the dynamic high link-error rate MANET with the high bandwidth but high delay satellite links creates a need to look at mobility and application QoS requirements in terms of routing in this environment. This work provides some insight into the impact of using OSPF in this environment as opposed to a MANET specific routing protocol.

2.3 Traffic Characterisation

The optimisation goal of traffic engineering (TE) is to support conflicting information exchange requirements while making the most efficient use of the currently available communication capabilities. Previous anecdotal evidence has suggested typical application types in this environment are text messaging, email, video, imagery, web, targeting, intelligence, collaborative planning, and voice. In order to understand the traffic characteristics of this type of network and what impact TE schemes may have, an accurate traffic model is critical. Figure 2 provides a more complete description of the traffic as seen during a Canadian naval exercise (Sibbald, 2004).

Fig. 2. Traffic Breakdown for a Naval Exercise

This chart divides the inbound traffic to a maritime node by application type. Internet and defence-based intranet web traffic took up a strong majority (91%) of the bandwidth. The remaining traffic was split between network overhead (i.e. routing), personnel and logistics management (PeopleSoft, MIMS, LotusNotes, SAP), email/collaboration (MS-Exchange, Outlook), and voice calls (VOIP). We will be using this mix in our traffic models described in Section 3.3.

Application/Network	Max Avg bps in/out	Peak bps in/out	Type	Priority
MCOIN (command and control)	24 / 35 K	45 / 80 K	Op	6
VOIP	5 / 16 K	50 / 140 K	Op	6
RSVP (network overhead)	continuous	.08 K	Net	5
OSPF (network overhead)	continuous	.26 K	Net	5
IGMP (network overhead)	continuous	.05 K	Net	5
TFTP (server to server)	0.3 / 0.6 K	22 / 30 K	Net	5
MS-Exchange (email)	30 / 48 K	60 / 130 K	Adm	4
Lotus Notes (Domino Replication)	0.2 / 0.5 K	18 / 38 K	Adm	4
DCOM (Outlook)	1 / 4.6 K	30 / 92 K	Adm	4
SAP (server to server)	0.7 / 1.2 K	28 / 64 K	Adm	3
Supply Program (MIMS/CFSSU)	0.1 / 1 K	3 / 10 K	Adm	3
Pay System (CCPS)	0.6 / 0.9 K	8 / 15 K	Adm	3
Pers Admin System (PeopleSoft)	2 / 4 K	6 / 30 K	Adm	3
Intranet (web)	6 / 8 K	60 / 100 K	adm	3
PC Anywhere (NM tool)	1.2 / 2.4 K	21 / 82 K	Net	2
Internet (web)	37 / 48 K	60 / 150 K	Rec	2
WindowsMedia (music/video)	7 / 15 K	35 / 120 K	Rec	2
MPEG Video (recreational)	2 / 34 K	30 / 64 K	Rec	2

Table 2. Application Bandwidth Requirements

During the same naval exercise, Table 2 was developed to specify fleet and ship data traffic usage and priorities. The maximum average and peak usage requirements for each of a variety of different types of applications as they are used in the navy are also provided. The traffic types shown are operational (Op), network overhead (Net), administrative (Adm), and recreational (Rec). One application that may not be obvious is Maritime Command Operational Information Network (MCOIN), which is the Canadian Navy's shore-based Command, Control, and Information System. The priorities listed provide an idea of the importance attached to the information being carried, and informs the network operator of how different traffic classes can be constructed to give preferential treatment within the network. Though the traffic mix and traffic priorities may vary over time it provides a reasonable description of the traffic that can be found in maritime networks, and more importantly the relative perceived priority of various traffic types. The priority information was used in the development of the traffic prioritisation service described in Section 4.2.

When multiple traffic types converge onto a single network, there is a requirement to ensure that time-sensitive (prioritised) information is delivered before less urgent traffic. Therefore, TE and communications management techniques must be applied to ensure that the priorities for information delivery are met. Our network services are described in Section 4.

3. Maritime Network Model

Based on this description of the maritime environment, a network model was developed using the OPNET discrete event simulation (DES) tool. In order to access the operation of the network services, several areas of the model had to be investigated. First, to determine the effect of network size, two network sizes were chosen based on maritime deployments; a small network, consisting of a NOC and a single four ship task force, and one larger network, consisting of a NOC and two four ship task forces. Second, the mobility of the two networks sizes was investigated. Finally, a maritime traffic model was developed.

3.1 Network Topology

Two network topologies were used in the simulations. The small network consists of five nodes and the large network has nine nodes as shown in Figures 3 and 4 respectively.

Fig. 3. Small Network

The connectivity of the small network model showing all the wireless links is shown in Figure 3. The link types are as follows. Ships 1-3 have satellite communications to the NOC (indirectly via satellite) with ship 2 and ship 3 using a 64kbps link while ship 1 has a 128kbps link. Each ship also has two 64 kbps radio links which form a ring. This implies that Ship 4 is only connected via LOS links from Ship 1 and Ship 3. The small configuration was designed based on the description of a single naval task force (AUSCANZUKUS, 2003).

Fig. 4. Large Network

The configuration of the large network model is shown in Figure 4. In this network Ships 1-3 and 5-7 have satellite links to the NOC with Ships 2, 3, and 7 at 64 kbps and Ships 1, 5 and 6 at 128 kbps. The line of site (LOS) links has been configured similarly for both four ship tasks groups. The large network configuration provides the opportunity to investigate more complex interactions between two task groups initially at some distance from each other. The mobility model described in Section 3.3 has the two task groups travel within LOS range of each other causing new connectivity and interference/bandwidth sharing.

In order to realise this network in OPNET, the base Cisco 7204 model was used for simulating the routing capability of the NOC. Ships use a custom built node model that includes capabilities for both point-to-point and wireless 802.11 links.

The point-to-point link model was used for satellite links as this most closely follows the leased bandwidth operation of satellite communications. The 802.11 link model was used for the LOS wireless links because it provides a wireless MAC that can simulate features such as fading and interference. The 802.11 model was modified to operate at the 64kbps LOS bandwidth rate and simulations indicate an operational throughput of approximately 42 kbps. One drawback of this approach is that while 802.11 uses CDMA, maritime LOS is most often TDMA. More work is required to validate our assumption that this difference is not significant at low bandwidth.

3.2 Mobility Model

The base geographical configuration of a task force is shown in Figure 5. With a LOS range of 18 Nautical Miles (Nm) and satellite capability for 3 ships the static topology on the small network shown in Figure 3 is achieved. The large network is composed of two task groups similarly configured but initially outside of LOS communication range of each other.

Maritime mobility is modeled here as a combination of two parts. Intra-task group mobility is based on the Nomadic Community model (Sanchez & Manzoni, 2005). Using this model the individual nodes of each task group move randomly within 3 Nm of their "base" position (as shown in Figure 5) causing links to fail when they exceed 18 Nm and recover when they are at most 18 Nm apart. Based on the nominal 30 nm/hour speed of maritime nodes, analysis suggests that LOS links in this model have a mean time between failures (MTBF) of about 5.5 hours and a mean time to recovery (MTTR) of 12.5 minutes. Note that

For the results described here, a single set of angles was used to simplify the simulations. The angle of arrival was set to 45° and the angle of trajectory was likewise set to -30° giving a trajectory similar to that shown in Figure 6. To give an idea of the impact of this aspect of the model, we discuss here the connectivity of ships 4 and 8. During the simulation, ship 4 is within range of ships 6 and 7 for an average of 66 minutes over the 130 minute run. Similarly, ship 8 comes within range of ship 1 for an average of 45 minutes during the simulation run. Note that though ships 4 and 8 do not have satellite communications, during the inter-task group mobility they come within range of ships in the other task group with high speed satellite (ships 1 and 6) and will thus prefer them over low speed satellite within their group. A 130 minute simulation time was chosen since it is also the time during which ships from one task group are within LOS range of each other

3.3 Traffic Model

Two different traffic loads were developed; nominal and heavy, with traffic distributions as given below. The resulting bandwidth, base OPNET traffic type, and application priority are given in Table 3.

Application	Avg bandwidth (kbps) nominal	Avg bandwidth (kbps) high	Type	Priority
Mcoin	.27 +/- .04 in .18 +/- .03 out	2.95 +/- .09 in 2.18 +/- .09 out	ftp	4
Voice Call	1.82 +/- .17 in 1.84 +/- .19 out	1.79 +/- .15 in 1.88 +/- .17 out	G. 729A	<=4
Over-head	.57 +/- .02 in .56 +/- .01 out	.56 +/- .02 in .57 +/- .01 out	ftp	3
Admin	.77 +/- .05 in .53 +/- .03 out	1.85 +/- .08 in 1.31 +/- .04 out	data- base	2
Intranet	10.57 +/- .41 in .65 +/- .02 out	11.93 +/- .53 in 0.83 +/- .05 out	http	1
Email	.45 +/- .06 in .39 +/- .05 out	1.06 +/- .09 in .81 +/- .08 out	smtp	1
Internet	14.97 +/- .56 in .97 +/- .05 out	22.07 +/- .60 in 2.71 +/- .06 out	http	0
Music/ Video	.30 +/- .03 in .13 +/- .05 out	.65 +/- .06 in .23 +/- .06 out	ftp	0

Table 3. Simulated Baseline Traffic

The nominal traffic models have been designed as closely as possible to the background traffic in maritime networks described in Figure 2 and Table 2. The traffic types in Table 2 have been simplified, with Overhead as an amalgamation of RSVP, OSPF, IGMP, and TFTP. Similarly, the Admin class encompasses Lotus Notes, DCOM, SAP, Supply Program, Pay system, and Personnel Admin System.

Based on the application, a corresponding OPNET traffic type was chosen and configured with an appropriate load. The bandwidth measurements shown in Table 3 were taken in

OPNET with all traffic (except voice) passing simultaneously across a 64kbps LOS link. This measurement was used to provide nominal traffic upon which our network services interacted. The voice call was measured separately over the same link. All bandwidths assume a normal distribution from 20 measurements with the given mean and a 95% two-way confidence interval. All measurements in this chapter are reported in the same way. The priority given in Table 3 corresponds to the priority given in Table 2 and is used to determine weightings in the traffic prioritisation service.

Traffic has been modeled based on pre-existing OPNET types as noted in the table. Also included is the priority of the application taken from Table 2. The high load traffic described above assumes increased traffic at times of high activity within the maritime network.

The high load traffic described above assumes increased traffic at times of high activity within the maritime network. Traffic has been modeled based on pre-existing OPNET types as noted in the table. The priority given in Table 3 corresponds to the priority from the network exercise in Table 2. QoS marking and associated WFQ weights are given in the TPS Section 4.2 below.

In both network topologies, traffic servers are on the NOC. This affects all traffic except for network overhead and voice calls. Overhead traffic is evenly spread between all nodes (including nodes in the other task group for the large network). Voice traffic is point to point and used only as noted for particular measurements.

4. Network Service Models

4.1 Traffic Monitoring Service

A Traffic Monitoring Service (TMS) was designed to measure the incoming and outgoing traffic of a node and distribute this information in summary form to all interested (subscribed) nodes in the maritime network. Currently, there is little if any logging of network traffic in maritime networks, and such monitoring is a critical need in order to provide the NOC staff with an up-to-date view of the operational state of the network. This is commonly termed the Network Common Operational Picture (NetCOP), of which this service would provide a part.

The TMS provides traffic information at a policy defined basis. Three “levels of detail” are supported; base, enhanced, and detailed. These levels provide increasingly detailed reports on the current traffic situation in the network but require increasing levels of bandwidth consumption and longer delays. The level of detail active at a particular time is tailored to the locally perceived load of the network. It is expected that base-level detail can be sent regularly without impacting network operations. Enhanced and detailed information can be sent intermittently or periodically at a very low rate. Timers and retransmissions of the summary data are used for fault tolerance.

In **base mode**, each node provides a summary of the aggregate traffic going in and out of that node (as total bandwidth or a percentage per traffic type). In **enhanced mode**, the service provides detail on the various priority levels and the current load at each level (tied to TPS service). More bandwidth is required to broadcast/multicast to all peer nodes in network. In **detailed mode**, traffic is further subdivided into individual long term flows and includes information such as delay, jitter and packet loss ratios (if available from protocols such as RTP). This is again node/link centered. Significant bandwidth is required to broadcast/multicast this information to all peer nodes in the network.

In order to model this service, three custom application types were created in OPNET, one for each level of detail. Loadings based on common SNMP style communications implied by the type of detail required were configured at each level. The impact of the different levels of traffic on the network and the delay and bandwidth requirements of each could then be studied. In Sections 5.2 and 5.3 we discuss the delay of the service and use the changes in the services delay to illustrate the utility of the traffic prioritisation and adaptive routing services described below. We also discuss the utility of switching between different levels of detail based on the current load on the underlying network.

4.2 Traffic Prioritisation Service

The Traffic Prioritisation Service (TPS) provides a mechanism to rank traffic by importance and prioritise resource allocation accordingly. It associates traffic to different classes of service that have relative priority between each another, and thus supports different forwarding requirements. Effectively, the service provides end-to-end (network-wide as opposed to a point-to-point) preferential treatment for certain applications. This allows relative traffic priority to be maintained from source to destination, including over the relay points. This preferential treatment is commonly known as DiffServ or soft QoS.

There are currently six classes of service: priority 0 (Best Effort), priority 1 (Background), priority 2 (Standard), priority 3 (Excellent Effort), priority 4 (Streaming), and priority 5 (Reserved). Weighted Fair Queuing (WFQ) was used, with WRED in the priority 0 (Best Effort) class. Resource allocations are given in Table 4 below.

Priority	Class Name	Weight	Notes
0	Best Effort	6	Recreational traffic
1	Background	6	Low priority applications
2	Standard	8	Operational applications
3	Excellent Effort	12	Routing and Management traffic
4	Streaming	18	Multimedia applications
5	Reserved	50	Up to 50% of bandwidth can be reserved for RSS flows.

Table 4. WFQ weightings for QoS used in measurements

In WFQ the relative weights correspond to the relative percentage of bandwidth that is assigned to each class of traffic. Since the weights assigned were engineered to add to 100, the assigned weight is the percentage of available bandwidth for each class if the link is fully loaded. Note that this means that if, for example, only one flow is in the standard class and there are three flows in the excellent effort class, the standard class flow will get at most 8% of the available bandwidth while each excellent effort flow will get an average of 4%. Thus bandwidth is assigned per class and not per flow. One of the most useful aspects of this scheme is what happens when one class is not fully saturated. Any bandwidth not used by a certain class is divided between the remaining classes, again in weighted order. Thus the reserved class (an additional priority class added to enable the RRS service described below) gains 50% of the bandwidth allocation not used by the other classes.

This service was one of the simplest to model in OPNET as it simply requires the application of existing QoS features. One of the interesting issues with OPNET was applying DiffServ to wireless models as this requires some understanding of the operational bandwidth in order

for the described weightings to be allocated correctly to the link. Since the bandwidth available on the link changes depending on environmental conditions and the number of other nodes transmitting on the same frequency, calculating the operational bandwidth requires extensive knowledge of the current state of the network. The model used here assumes a nominal bandwidth of 42 kbps over a two-user 64 kbps LOS link.

4.3 Adaptive Routing Service

Since maritime nodes may have multiple WAN links of varying capacity, it may be useful for applications to use only a subset of the available links. This may be for reasons of delay sensitivity (e.g. for VOIP calls) or because of the bandwidth capabilities and error/failure rate of the link (e.g. for ftp communications).

The Adaptive Routing Service (ARS) provides an alternative routing method for matching a traffic class to WAN resources. Essentially, it indicates what types of traffic must/should travel over a certain type of bearer. It makes use of resource availability (i.e. does the bearer possess sufficient bandwidth to meet the requirements of that traffic class) and resource suitability knowledge (i.e. will the bearer meet the QoS requirements of that traffic class).

The benefits of such routing flexibility are significant. Besides ensuring that traffic of a given class will flow over a bearer that supports it, ARS offers a solution to the well-known load balancing problem. Traffic from the same source and destination can be directed to travel over different routes. Since the path selected is also based on the traffic type instead of simply the route cost (shortest path), ARS provides a better distribution of traffic across the WAN links and thus better utilisation of available network bandwidth.

The ARS was modelled through the use of MPLS tunnels. An overlay on the network was created specifically for VOIP calls. VOIP calls were then routed over the least loaded links. In an actual implementation, multiple MPLS overlays could be created to avoid links that cannot support the QoS requirements of the application flow.

4.4 Resource Reservation Service

The RRS has been designed with a number of features to deal with the specific requirements of maritime environments. This includes probing multiple routes in parallel for load balancing and an increased acceptance rate, a priority and pre-emption scheme to favour the most critical flows, fault tolerance mechanisms to deal with mobility and link errors, and dynamic reconfiguration of its parameters to meet changes in operational requirements.

The RSS consists of four phases. In the first phase, global link information is used to generate multiple routes between the source of the requesting flow and the destination. The second phase of the algorithm probes the potential routes separately to determine if sufficient resources are available (which increases the acceptance rate). In the third phase, an acceptable path is selected and committed (which promotes load balancing), potentially pre-empting lower priority flows (which prioritises the most critical flows). Finally, in the fourth phase, the reservation is maintained until the traffic flow is terminated by the user, the reservation lifetime ends, the reservation is ended manually, or the network can no longer support its requirements. Mobility is handled by assuming the network is stable for the period of call setup, and network maintenance handles topology changes while the reservation is active as described below. A detailed functional description of this service is given in (Kidston et al., 2007)

4.4.1 Reservation Protocol Models

In order to evaluate the impact of the specialised features of the RRS in maritime networks, our simulations compare it with both RSVP, a standard reservation protocol for fixed networks, and INSIGNIA, a reservation protocol proposed for MANETs. There are significant differences in the operation of the three reservation protocols simulated. The largest difference is that RRS includes multi-routing and pre-emption. In RRS, each reservation is made with up to 3 parallel probes to exercise the partially disjoint multi-routing aspect of the protocol. The priority mechanism of the service was also exercised by assigning each new reservation one of three priority levels (low, medium, and high) with equal probability. Requests of lower priority may be pre-empted (dropped) in order to admit a higher-priority flow. High-priority flows are not pre-empted, and may only be blocked from being accepted by other high priority flows. Pre-empted flows are called admitted (they were initially accepted), but unsuccessful (they terminated before their scheduled end time). If a request is not initially admitted, it is also unsuccessful. No additional attempts are made to establish a reservation.

In order to implement the RRS, some OPNET models had to be modified and new ones created. Our approach was to use the existing IP networking models and simply add RRS packet processing capability on top, so RRS messages could be processed and forwarded as required. First, a significant change was required to the OSPF model in the ship's routers. In order to capture changes in network topology and determine the type and nominal bandwidth of the link from the OSPF cost, a software tap was added to the existing OPNET model. This tap forwards all link state database (LSDB) changes to the local RRS model, which maintains its own internal representation of the network connectivity. Two additions were made to the network node models for simulating RRS. The first was a simple process for generating reservations that submits new request interrupts at a configurable rate to the RRS process. The second was the RRS process itself, which includes all the logic previously described for forwarding packets and reserving resources.

The RSVP model was based on the description in (Braden et al, 1997). The main differences from the RRS model is the lack of multi-routing (only the default route is probed), pre-emption (no prioritisation method is included in RSVP), and fault tolerant features (timers and retransmissions of RRS packets).

The implementation of the INSIGNIA (Lee et al., 2001) model was also derived from the RRS model, though in this case significant changes were required. Since INSIGNIA reserves resources per hop, resources are reserved right away if available, otherwise no resources are reserved further in the route and a report is sent from the destination that the request was not successful. If resources are reserved all the way to the destination, success is reported.

4.4.2 Reservation Request Models

In order to assess the operation of the reservation protocols, two variables were investigated. First, the impact of the **source** of requests was investigated with two different models. Second, in order to determine the effect of network **loading**, two network reservation request arrival rates were chosen.

For all protocols, the total time a reservation remained active was based on an exponential distribution with mean of 270 seconds. Reservations are for 8 kbps, with a maximum of 50% of each link's bandwidth available for reservations. These values were chosen to simulate voice connections.

The source of reservations arriving in the network was varied to investigate the impact of the multi-routing aspect of the RRS. Reservations may either originate uniformly from all nodes in the network (uniform model), or originate only from a single node (single source model). In the uniform model, the request generation process was activated on all nodes, while in the single source model the request generation process was activated only on a single node chosen randomly at the beginning of each simulation.

Considering the effect of different request source models, four reservation inter-arrival rates were used to simulate reservation saturation (nominal loading) and reservation overload (high loading). The inter-arrival time for reservations using the uniform model were exponentially distributed and centered on 60 seconds for nominal request load and 30 seconds for high load. These rates were chosen to saturate and overload the network with reservations respectively. For the single source model, the inter-arrival time was set to 30 seconds for nominal load and 15 seconds for high load for the same reasons. Note that since the request loads are not the same for uniform and single source request models, the results of these two source models should not be directly compared.

5. Results

5.1 Methodology (Credibility)

Wireless network simulations research has come under increasing scrutiny of late. Recent publications have raised questions about the credibility of past simulation work and have suggested different methods by which this can be resolved. For example in (Andel & Yasinsac, 2006) the lack of credible results in mobile ad-hoc network (MANET) simulations are blamed on a number of systemic problems. Without documenting all settings and data sets, simulations are not repeatable. Without addressing the sources of randomness and the data collection techniques, simulations cannot be statistically valid. Without comparing results with a real-world implementation, simulations cannot be empirically validated. Finally, without identifying the scenario, the traffic will not be complete (i.e. unrealistic). Similarly, in (Kurkowski et al, 2005) the authors studied a collection of 114 peer-reviewed MANET simulation papers presented at the MobiHoc symposium between 2000 and 2005. They found that 85 percent of the papers were not rigorous because they did not specify all parameters used in their simulations. For example, 30 percent of the papers did not identify the simulation environment used. Others did not include parameters such as transmission range, number of simulation runs, traffic type or mobility model. They focused on the need for unbiased and statistically valid methods.

As an exercise we decided to investigate what it would take to make our own simulation results properly credible. For this work we have combined the issues into a single list and attempted to apply these five principles to our own simulations. In order for simulation research to be credible it must be:

- **Repeatable**, where experiments describe all configuration settings;
- **Rigorous**, where the model settings varied, and how much they are varied, exercise the feature under investigation
- **Complete**, where the model is not oversimplified (avoiding ambiguous or incorrect conclusions)
- **Statistically Valid**, where the method of analysis is described and follows mathematical principles;

- **Empirically Valid**, where simulations are compared against a real world example.

To provide **repeatability**, the following settings were used in all simulations except where specifically noted. OPNET version 11.0 PL1 was used with the node and link models as noted below. Additional models were created or existing models were modified in some cases in order to model the RRS. Runs of 130 minutes were used for all measurements. Statistics gathering began after 270 seconds. This value was chosen because it is approximately three times the amount of time required for routing to converge and applications to reach steady state. The OSPF routing protocol was used with a hello interval of 10s, dead interval of 40s, delay of 1s, and retransmission interval of 5s.

In order to provide **rigorous** and **complete** results the following approach has been taken. The main metrics of interest in the work relate to the acceptance and pre-emption rate engendered by RRS and how those rates are affected by various request loads and source distributions. The simulation setup described below provides the complete set of network level configurations that were changed to provide minimal but sufficient variability to exercise the RRS as described here. The variability in the application performance can thus be fully ascribed to the changes in request configuration and the RRS itself.

To ensure the results described here are **statistically valid**, the following approach was taken. Twenty runs were made for the simulations to have tight error bounds. Statistics are averaged over each run. All results are quoted with a 95% confidence interval, which gives values within the specified range 19 times out of 20. The mean is calculated by summing the result of each run and dividing by the number of results. The standard deviation is calculated as the square root of the variance of this mean. From this the standard error is calculated as the standard deviation over the square root of the number of results. Finally, the two way 95% confidence interval is calculated as an (+/-) offset of the mean with a value 2.093 times the standard error for 20 measurements.

For **empirical validity**, the simulations were previously compared with an existing prototype that implements RRS in a wired test-bed (Kidston & Kunz, 2008). In that paper, we showed the RRS simulation results matched the operation of a prototype.

5.2 TMS Simulation results:

Based on the topology, mobility and background traffic described in Section 3, the delay of injecting Traffic Monitoring Service (TMS) traffic into the network was measured as shown in Table 5 and Table 6.

	Nom Load delay (s)	High Load delay (s)
Base Mode	3.8 +/- 0.5	6.9 +/- 0.7
Enhanced Mode	13.2 +/- 0.9	23.6 +/- 1.2
Detailed Mode	27.8 +/- 2.2	55.7 +/- 2.5

Table 5. TMS Delay in seconds, Small Network

	Nom Load delay (s)	High Load delay (s)
Base Mode	4.3 +/- 0.3	7.2 +/- 0.6
Enhanced Mode	15.6 +/- 0.9	28.2 +/- 2.2
Detailed Mode	35.2 +/- 1.8	61.5 +/- 3.2

Table 6. TMS Delay in seconds, Large Network

As can be seen, the effect of increased load is readily apparent in maritime networks, with the TMS delay almost doubling from nominal to heavy background traffic. For both the small and large network the base mode delay during nominal load is approximately four seconds, which for a non critical network service is most likely acceptable. However the enhanced and detailed modes have a much longer delay. This may be acceptable if the information is not being used interactively.

In order to investigate the impact of adding adaptability to this process, the service was modified to switch between detail modes to limit the maximum delay while delivering the most information possible. The following graph (Figure 7) shows the effect of adaptability on the operation of the TMS. Note that in this graph the TPS service described in Section 5.3 was also active.

In this simulation the TMS is attempting to ensure that the response time is at most 30s by reducing the detail mode to base if the response time exceeds 60s and to enhanced if it exceeds 30s. Similarly it will increase the monitoring to enhanced if it is less than 3s and currently at base and to detailed if it is less than 3s and currently at enhanced. In Figure 7 the small network begins with no background traffic. At 20 minutes, heavy background traffic is added. Nominal background traffic began at forty minutes. The figure shows clearly the switch from detailed mode after a 90+ second delay after 20 minutes and to enhanced mode after a 2 second delay after 40 minutes. Different policies of when and what should cause the switch between detail modes could also be implemented. This result shows that the causes and results of such changes can be modelled and compared in OPNET.

Fig. 7. Traffic monitoring adaptability example - small network - dynamic load

5.3 TPS Simulation results

The TPS service described in Section 4.2 was enabled and the TMS simulations rerun with the following results.

(with QoS)	Nom Load delay (s)	High Load delay(s)
Base Mode	2.3 +/- 0.3	4.6 +/- 0.5
Enhanced Mode	10.4 +/- 1.3	15.0 +/- 1.4
Detailed Mode	22.2 +/- 2.1	37.2 +/- 2.5

Table 7. TMS Delay in seconds, Small Network with TPS

(with QoS)	Nom Load delay (s)	High Load delay(s)
Base Mode	3.4 +/- 0.3	5.7 +/- 0.7
Enhanced Mode	12.3 +/- 1.1	20.3 +/- 2.3
Detailed Mode	26.8 +/- 2.0	45.2 +/- 3.2

Table 8. TMS Delay in seconds, Large Network with TPS

Table 7 and Table 8, when compared with Table 5 and Table 6 respectively, show a significant improvement in the TMS delay. This confirms that DiffServ-style QoS can reduce the delay of prioritized flows in this environment. Further studies in OPNET could be completed to determine the impact of alternative WFQ weightings and alternate operational bandwidth assignments to the LOS links.

In another test to gauge the effect of TPS on different types of traffic, the delay of a voice call between the NOC and ship 4 in the small network was measured with and without TPS enabled. The default route for such traffic is to relay through ship 1. At low load one call was made with priority 4 and the delay measured was .67 +/- .12 seconds without TPS and .13 +/- .01 with. At high load, two identical calls were made; one at priority 2 and the other at priority 4. Without TPS, the end to end packet delay was the same for both calls at 1.4 +/- 0.3 seconds. With TPS, the high-priority call had a delay of 0.7 +/- 0.2 seconds while the low-priority flow's delay was 1.4 +/- 0.4 seconds. Since an acceptable voice delay is approximately 500ms TPS enables a single acceptable voice call at nominal load, but at high load two voice calls are not possible even with TPS. Note however that since background traffic uses the default route via Ship 1, the other LOS link via Ship 3 is currently unloaded.

5.4 ARS Simulation results

In order to improve utilization of the network, an MPLS overlay was introduced to allow traffic travelling from the NOC to Ship 4 to take different routes depending on the application type and priority. In this case high priority voice traffic was to travel via Ship 3 while all other traffic will travel over the default route via Ship 1. When this was done, the load on the Ship 1 to Ship 4 LOS link was reduced from an average utilization of 90.5% to 10.8% while the loads on the alternate LOS link from ship 3 to ship 4 was increased from almost nothing to 10-16%. The large reduction in average bandwidth on the default route is partly caused by a reduction in TCP based retransmissions.

With the combination of TPS and ARS, the impact on the delay of voice packets is significant. The high-priority voice call taking the alternate lightly loaded route via Ship 3 has a delay of 0.19 +/- 0.03 seconds while the lower priority voice call with the default route

has a delay of 0.43 +/- 0.10 seconds. This arrangement using the combination of ARS and TPS made the high priority flow of acceptable and the low priority flow at least marginal.

5.5 RRS Simulation results

This section describes the RRS results. Results for RSVP and INSIGNIA are included for comparison. First, the acceptance rates of RRS and RSVP are compared in several parts. Next, the different nature of maintaining reservations in INSIGNIA leads to an alternative comparison. Finally some conclusions on the performance of RRS are given.

5.5.1 RRS vs. RSVP, Static Network Model

Our evaluation begins with the acceptance rates in RRS and RSVP in a static network (no mobility). Table 9 provides the percentage (%) of the requests that were able to reserve resources from source to destination at the time of the request. A margin of error is given at the 95% confidence interval.

Network	Load	Source	RRS	RSVP
Small	Nominal	Uniform	93.1 +/- 0.6	78.1 +/- 1.0
		Single	91.2 +/- 0.7	64.9 +/- 1.2
	High	Uniform	75.3 +/- 1.0	57.3 +/- 0.7
		Single	67.6 +/- 1.1	40.2 +/- 0.8
Large	Nominal	Uniform	88.8 +/- 0.5	67.8 +/- 0.8
		Single	88.4 +/- 1.1	58.4 +/- 1.5
	High	Uniform	68.6 +/- 0.7	48.6 +/- 0.5
		Single	65.2 +/- 1.0	37.3 +/- 0.7

Table 9. Acceptance Rates, Static Network

The most immediate conclusion that can be drawn from Table 9 is that RRS provides superior acceptance rates to RSVP in all scenarios. An improvement of 19-41% over RSVP is achieved when the source of requests is uniformly distributed, and an improvement of 41-75% with a unique source of reservations. However, the reservation success rate, defined as a reservation which gains end-to-end resources from the beginning to the end of its request, should also be considered. In this case the reservations lost to pre-emption in RRS reported in Table 12 must be included. Since these protocols use the same two-phase commit strategy for reserving resources, the improvement by RRS can be attributed primarily to two factors: the use of pre-emption to admit higher priority flows; and the use of multi-routing to route around congested links. These effects are discussed in more detail in Section 5.5.3.

5.5.2 Effect of Mobility on RRS and RSVP

The acceptance rates of RRS and RSVP were also simulated using the mobile network model, with results shown in Table 10.

A comparison of Table 9 and Table 10 shows that the mean acceptance rates of the mobile network are generally lower than in the static case, i.e. within or below the 95% confidence interval of each other in all but two cases. In the large network at nominal load, the single source model of RRS has a mean in the mobile case 3.3% above the mean of the static

network while at high load RRS similarly has a mean 1.3% above the mean of the static network using the single source reservation model. This would suggest that mobility has a small negative effect on raw acceptance rate in the small network, with a more variable effect in the large network.

Network	Load	Source	RRS	RSVP
Small	Nominal	Uniform	91.6 +/- 0.7	77.0 +/- 0.9
		Single	90.5 +/- 1.0	63.5 +/- 1.7
	High	Uniform	74.1 +/- 0.8	56.5 +/- 0.9
		Single	67.2 +/- 1.4	40.1 +/- 0.7
Large	Nominal	Uniform	88.1 +/- 1.1	64.7 +/- 1.0
		Single	91.7 +/- 1.1	57.5 +/- 1.3
	High	Uniform	67.5 +/- 0.5	48.3 +/- 0.6
		Single	66.5 +/- 1.4	37.7 +/- 1.0

Table 10. Acceptance Rates, Mobile Network

The effect of link failures on active reservations is related in Table 11 with the given percentage of accepted flows having lost their resources at some point along their route.

Network	Load	Source	RRS	RSVP
Small	Nominal	Uniform	2.0 +/- 0.6	1.7 +/- 0.6
		Single	1.5 +/- 0.7	1.7 +/- 0.8
	High	Uniform	1.3 +/- 0.6	1.6 +/- 0.5
		Single	1.0 +/- 0.3	1.0 +/- 0.7
Large	Nominal	Uniform	4.3 +/- 0.3	4.3 +/- 0.3
		Single	4.8 +/- 0.9	4.2 +/- 0.9
	High	Uniform	3.6 +/- 0.3	3.8 +/- 0.4
		Single	3.7 +/- 0.9	4.4 +/- 0.8

Table 11. Reservation Failure Rates (due to mobility)

The direct comparison of Table 10 does not take into account the reservations later lost to the link failures associated with mobility. Mobility can cause existing successful reservations to be lost when links fail, thus increasing the number of subsequent reservations admitted as shown in Table 11.

The mean failure rates for RRS and RSVP can be seen to fall within the 95% confidence interval of each other in all cases. This is as expected, since they are based on the same underlying mobility model and a similar reservation release mechanism. Reservation recovery mechanisms were not included in the RRS and RSVP model. Considering the relatively low number of failed flows relative to the number of accepted flows, it is unlikely that such features are worth the additional overhead in this low bandwidth environment.

Considering the link failure rate, the total number of successful reservations can be calculated to determine the effect of mobility on RRS and RSVP individually. A successful reservation is defined as a reservation that maintains their resources end-to-end without loss due to a link failure or pre-emption. In this section we look only at link failures. In RRS, the effect of mobility (failure rate) with the single source reservation model has very little effect, with only 1.2-2.3% fewer successful reservations with mobility when compared to the static

case. The effect is slightly larger in the uniform source model with 2.9-5.1% fewer successful reservations overall. RSVP shows a similar trend, though with a slightly larger effect. Compared with the static model, 1.2-5.7% fewer reservations were successful in the mobile network for single sourced reservations, while the uniform model had 3.0-8.7% fewer with mobility. The difference between single source and uniform models is explained by the fact that the uniform model saturates the links more evenly, while the single source model suffers from bottlenecks around the reservation source. This leads to more reservations on average being lost for a particular link failure in the uniform model. From this we conclude that there is a slight (single digit percent) negative effect from mobility on reservation success, with uniform reservations experiencing approximately double the effect found using the single source request generation model. In order to properly compare RRS and RSVP using the idea of successful reservations, we look in the following section at the other cause of reservation failures, pre-emption.

5.5.3 Effect of Pre-Emption

By investigating the effect of pre-emption rates in RRS, we gain a better understanding of the difference in reservation success between RRS and RSVP. The percentage of accepted flows which lost their resources due to pre-emption is given in Table 12.

Network	Load	Source	RRS (static)	RRS (mobile)
Small	Nominal	Uniform	8.3 +/- 0.5	8.2 +/- 0.7
		Single	15.0 +/- 1.1	15.8 +/- 1.3
	High	Uniform	21.7 +/- 0.7	21.0 +/- 1.2
		Single	35.8 +/- 0.9	35.7 +/- 1.2
Large	Nominal	Uniform	8.9 +/- 0.5	8.7 +/- 0.4
		Single	18.1 +/- 1.1	17.1 +/- 1.0
	High	Uniform	19.2 +/- 0.5	18.9 +/- 0.5
		Single	34.4 +/- 0.8	34.2 +/- 1.0

Table 12. Pre-emption Rates (RRS only)

From this table we can see that pre-emption is significantly impacting existing RRS flows, particularly in the high load scenarios. In the maritime environment, this level of loss may be acceptable considering that no high-priority flows are affected, only low priority and to a lesser extent medium priority flows. This ensures the acceptance of high-priority reservations, except in extreme cases, where they may be blocked by other high-priority reservations. This is unlikely to occur in even the high-load models simulated here given the relatively low pre-emption rates and an even distribution of requests between the three priority classes.

Comparing the static and mobile network model results, the pre-emption rates are within the 95% confidence interval of each other in both the static and mobile scenarios. This is to be expected, since with similar reservation rates in both mobile and static case, the mix of reservations in the network is similar. With similar network priorities and similar number of reservations, the pre-emption rate should also be similar. Though within error bounds, a slightly higher pre-emption rate in the static case can be seen. Since there are slightly more reservations made in this case, additional pre-emption can be expected.

In order to quantify the effect of priority on acceptance and pre-emption rates in RRS, we investigated the large static network scenario with uniform high traffic. Priority was found to have a significant impact on acceptance rate, with high priority traffic having an acceptance rate of 87.9 +/- 0.7 percent while medium and low-priority flows had an acceptance rate of 64.7 +/- 0.9 and 49.7 +/- 0.9 percent respectively, for an acceptance rate of 68.6 +/- 0.7 percent overall. Similarly, while high-priority flows were not pre-empted, medium-priority flows had a pre-emption rate of 25.4 +/- 0.7 percent and low-priority flows had a pre-emption rate of 42.8 +/- 1.6 percent, for a pre-emption rate of 19.2 +/- 0.5 percent overall. This shows that priority has a significant impact on both acceptance and pre-emption rates, with high-priority flows gaining service similar to RSVP (i.e. no pre-emptions) but with an improvement of 80.9 percent in mean acceptance rate over RSVP for the large static network scenario with high traffic.

The amount of pre-emption measured, especially at high load, gives rise to the question of whether RRS is in fact an improvement on RSVP in terms of successful reservations. Simple subtraction of the pre-emption rate from acceptance rate is however not appropriate, as reservations must have achieved their resources for at least some period of time in order to be pre-empted. Based on the percentage of accepted flows that were not pre-empted (or lost due to link failures) the reservation success (completion) rate improvement of RRS over RSVP can be measured. Analysis shows there is a large difference in mean improvement rates in high vs. nominal load scenarios. At high load, an improvement of only 3-8% more successful reservations over RSVP can be achieved in the small network and 14-17% in the large network, regardless of mobility or traffic source model. At nominal load, a greater improvement is possible, in the small network 9% and 20% for uniform and single source models respectively. In the large network at nominal load there are some mobility effects. The static network gains 19% and 24% for uniform and single source models respectively, and the large mobile network RRS reservations gain 24% and 31% for uniform and single source models respectively. This shows that at nominal load the multi-routing effect is especially effective for single sourced requests while at high load there is little difference between the two request models.

It should be noted that though RRS does pre-empt low-priority flows, these flows gain some advantage from the use of reserved resources for the period of time before they are pre-empted. Investigating the effect of priority level on resource hold times (reservation success) we again looked at the large static scenario with uniform high traffic. In this scenario we found that high-priority flows were not pre-empted (as expected), but both medium and low-priority flows which were accepted had on average a significant period in which they did gain their required resources. Medium-priority flows that were eventually pre-empted kept their reserved resources for 65.2 +/- 4.2 percent of their allocated time period on average. Similarly, low-priority flows maintained their reserved resources for 36.6 +/- 2.7 percent of their allocated time. Thus, though pre-empted flows do not gain full advantage of reserved resources throughout their lifetime, RRS does provide them with significant periods of advantage based on their priority.

5.5.4 RRS vs. INSIGNIA

In order to compare RRS with INSIGNIA, it is important to remember that in INSIGNIA flows are granted resources per-hop for as far along their current route as they are available instead of end-to-end. This means that if a link does not have resources, later links in the

flow will not reserve resources. The unfortunate consequence seen in these simulations is that resources are kept by flows on the first part of their path, and yet flows still fail to achieve end-to-end reservations. As shown in Figure 8, this reduces the total number of successful end-to-end reservations in the network because resources are wasted on non-viable reservations.

Fig. 8. Example of Network-Wide Reserved Bandwidth in INSIGNIA

Figure 8 shows the total amount of reserved bandwidth on all links in the network at a particular time in one simulation run of the large network with no mobility, nominal request load, and with the uniform reservation arrival model. It shows that at the beginning of the simulation, from about 5 minutes to 25 minutes, a large number of reservations are successful and the total amount of end-to-end reserved bandwidth in the network peaks around 1800 – 1900 kbps (out of a theoretical maximum of 2176 kbps “reservable”). After this, the total amount of reserved bandwidth decreases until a steady state is achieved at about 40 minutes. From this point on, successful reservations hold approximately 1250 kbps of network bandwidth. The remaining reservable bandwidth at this point is tied up by reservations which do not have end-to-end resources but are still holding resources on the beginning of their route, blocking other reservations from getting sufficient resources to be successful themselves.

In Table 13 and Table 14 below, the acceptance rate given is the percentage of new flows which gain resources on all links in their route on the first try. Upgrades are the percentage of flows which at some point did not have end-to-end resources then gain such resources. Downgrades are the percentage of accepted or upgraded flows which at some point had end-to-end resources and then lose the resources on any link. Since there is no pre-emption in INSIGNIA this can only happen because of mobility.

Network	Load	Source	Acceptance	Upgrade
Small	Nominal	Uniform	20.8 +/- 1.2	16.3 +/- 1.1
		Single	14.4 +/- 0.6	17.7 +/- 1.0
	High	Uniform	6.7 +/- 0.3	14.8 +/- 0.6
		Single	6.6 +/- 0.7	10.1 +/- 0.7
Large	Nominal	Uniform	22.7 +/- 0.9	6.9 +/- 0.5
		Single	47.5 +/- 2.1	5.3 +/- 0.8
	High	Uniform	9.7 +/- 0.4	8.7 +/- 0.4
		Single	27.9 +/- 0.8	4.7 +/- 0.5

Table 13. INSIGNIA Results (Static Network)

Network	Load	Source	Acceptance	Upgrade	Downgrade
Small	Nominal	Uniform	18.1 +/- 0.9	18.5 +/- 0.9	4.9 +/- 1.0
		Single	14.6 +/- 1.7	20.6 +/- 1.8	5.1 +/- 2.3
	High	Uniform	6.2 +/- 0.3	14.7 +/- 0.7	3.5 +/- 1.2
		Single	7.1 +/- 0.6	12.2 +/- 1.4	4.3 +/- 1.6
Large	Nominal	Uniform	19.7 +/- 0.8	10.4 +/- 0.6	8.8 +/- 1.1
		Single	45.1 +/- 1.6	11.5 +/- 1.6	4.4 +/- 1.2
	High	Uniform	8.5 +/- 0.3	9.3 +/- 0.5	11.3 +/- 1.1
		Single	27.3 +/- 1.0	9.7 +/- 0.8	7.5 +/- 0.8

Table 14. INSIGNIA Results (Mobile Network)

Considering the results of these two tables, it can be seen that INSIGNIA performs very poorly in the maritime environment with low acceptance rates (most below 30%). These results would not be acceptable in a maritime environment, especially considering the lack of priority mechanisms for critical flows.

Comparing these two tables to determine the effect of mobility, it can be seen that the static network model provides a slightly higher initial acceptance rate, which is to be expected when links may be down due to mobility when new requests arrive. Comparison of the acceptance rates show the static results are within 13 percent of the mobile results in all cases respectively. Conversely, upgrades are higher in the mobile network. This is due to the fact that when links become available due to mobility there is a greater chance for existing reservations to gain end-to-end resources using the newly available link. When both upgrades and downgrades are taken into account, the static and mobile results for partially successful reservations are similar and within +/- 17%. Partially successful reservations are defined as reservations which achieve end-to-end reservations at some point in their lifetime. Interestingly, the uniform source distribution resulted in 7-17% more partially successful reservation in the static network model (compared to the mobile model) while the single source distribution resulted in 2-10% less. Because fewer links become fully subscribed in the single source distribution due to the bottleneck around the source, it does not suffer as many lost reservations when a link fails as, on average, there are fewer reservations in the network. Uniformly distributed requests are conversely more sensitive to link outages since all links are more likely to have a high number of reservations.

5.5.5 Conclusions of RRS Results

Looking at our results in terms of the operational requirements for maritime networks, the overall RRS acceptance rate of 91-93% on average for nominal loadings regardless of mobility is acceptable. In the critical high load case, the RRS acceptance rate of 67-75% on average may seem low but it should be noted that the pre-emption mechanism used in RRS ensures that high-priority flows are accepted at the cost of lower-priority flows losing their resources. For example, in the heavily loaded static network with uniform requests, 87.9% of high priority traffic was accepted on average, while low-priority traffic was accepted only 49.7% of the time on average.

To evaluate the effectiveness of RRS in a maritime environment with dynamic topology, it was compared with the archetypical fixed network reservation protocol RSVP and a MANET reservation protocol INSIGNIA. With mean acceptance rates of 64-78% on average at nominal load and 40-57% at high load it is unlikely that RSVP would be acceptable in this environment. INSIGNIA's performance was even worse with mean acceptance rates of 14-21% at nominal load and 6-7% at high load. In a raw comparison of acceptance rates, RRS is 19-76% better than RSVP and 86-1095% better than INSIGNIA.

From these results, it can be seen that the multi-routing and pre-emption features of RRS provide a higher acceptance rate compared with RSVP with similar loss rates during link failures. This improved acceptance rate does however come at the cost of pre-empted lower priority flows. In order to determine the impact of pre-emption, RRS and RSVP were compared in terms of successful reservations which maintain their resources end-to-end throughout their lifetime. It was found that RRS still outperformed RSVP by 3-8% at high load and 9-20% at nominal load. These numbers highlight that probing multiple routes makes a significant difference only when the network is not already saturated with requests. Another interesting conclusion from these simulations is that the mobility models simulated have only a marginal negative effect on both acceptance rates and reservation success. Comparing the results for the different mobility models, the acceptance rates for RRS and RSVP with mobility are within or slightly below the 95% confidence interval of the static model in most cases.

6. Related Work

There has been, to the best of our knowledge, no network modelling work on maritime networks. However, there has been some recent work on improving networking in this area, for instance in applying static DiffServ QoS to maritime networks (Barsaleau & Tummala, 2004). This paper showed that throughput and delay guarantees were hard to achieve in this environment, but queuing and dropping mechanisms, if properly tuned, could provide limited service differentiation. This work does not consider the dynamic nature of the maritime environment, where the importance attached to different classes or even flows of information vary with time. The use of modelling in this environment would greatly aid investigations in the type of tuning required in different circumstances before incurring the expense of deployment on operational platforms.

The Resource Reservation Protocol (RSVP) (Braden et al, 1997) is a well known standard that reserves resources for unicast or multicast flows along the default path(s) from sender to receiver(s). RSVP delivers quality-of-service (QoS) requests to all nodes along the path(s) of the flows and establishes and maintains "soft" state related to the requested service. This

provides support for dynamic reservation membership and automatic adaptation to routing changes. During reservation setup, RSVP transports traffic and policy control parameters that provide direction to nodes as to whether the flow should be admitted. From the point of view of maritime networks, while RSVP does provide a basic end-to-end service, it uses the default route (which is quickly saturated), does not support prioritisation (a requirement), and assumes network reliability (lacks fault tolerance mechanisms). These issues were addressed in the RRS design.

The adaptation of existing fixed-network-based QoS mechanisms into a MANET environment has been investigated from many different angles, all of them distributed. INSIGNIA (Lee et al, 2001) is an IntServ-based in-band signalling system for providing QoS reservation services on top of existing MANET routing protocols. The INSIGNIA framework supports distributed resource reservation, restoration, and end-to-end adaptation irrespective of the underlying routing protocol. Reservations are accomplished through the use of a specialized IP option field added to all packets. When a packet arrives at a node, a reservation is made on the outgoing link as long as the reservation was successful so far on its route and there are sufficient resources on the local link. End-to-end adaptation to available bandwidth is possible through the use of user-supplied policies that inform applications as to the available bandwidth reported by the protocol. As was seen in the simulations, this single phase commit strategy fails in congested networks.

7. Conclusions and Future Work

In this chapter we have described a network model of maritime networks and, using modelling and simulation, investigated the impact of several network services designed to provide traffic engineering. In order to add credibility to our results we followed a process to ensure it would be: repeatable, by including all parameter settings; rigorous, where the variables investigated are well described and correctly exercise the simulated model; complete, in that the model is sufficiently detailed; statistically valid, in that the mathematical analysis is sufficient and correct; and empirically valid, such that the results of the simulation match well with a real world example.

Four network services were described. The traffic monitoring service (TMS) provides details on the traffic generated, received and passing through the local node. Three levels of detail are possible, and the amount of information sent to interested parties, most often simply the network operation centre, can be tuned so that the status information is delayed by a maximum amount. Simulations showed that in both small (four ship) and larger (eight ship) networks, the data could be updated at least every 30 seconds by switching to lower levels of detail when the intervening network was loaded.

The second type of service was the traffic prioritisation service (TPS). By assigning applications to different DiffServ classes, the delay of more important application traffic could be reduced. Simulations showed that the TMS delay could be significantly reduced when assigned a high priority. It also showed that in a small network voice calls that had unacceptable delay without TPS could be made acceptable with TPS.

The third network service, the adaptive routing service (ARS), uses MPLS to divert traffic from links that cannot meet the QoS requirements of the application. Simulations showed that alternate routes could be used to make voice calls with unacceptable delay

characteristics acceptable. This shows that altering the relationship between queuing resources and applications can achieve a desired service level for some traffic.

The fourth network service, the resource reservation service (RRS), uses distributed two phase admission control for guaranteed end-to-end bandwidth reservations. This service includes several novel features designed specifically for maritime networks including multi-path probing and bi-directional reservations. The value of multi-path probing is demonstrated by simulation giving RRS an acceptance rate 19-76% better than RSVP and 86-1095% better than INSIGNIA, two alternative reservation protocols.

Our main conclusion is that, though non-trivial, it is possible and valuable to create a maritime network model using OPNET. Such models can be used for testing the impact of new network applications and services as well as tuning existing services. It also showed that altering routing can lead to more optimal use of link b/w resources, and can also lead to better QoS for critical traffic. The combination of TPS and ARS provided improved delay, but there can still be problems for critical application flows in times of high usage, congestion, and low connectivity from mobility. For this case RRS provides a high probability resource reservation service (RRS).

Several simplifying assumptions were made in this work including the use of 802.11 models instead of true VHF LOS models. Further work is required to see if this was a valid replacement.

There remain many possible avenues to further this work on evaluating network services for the maritime environment. Other scenarios could be investigated with more ships, and different mobility models. There is also quite some scope to further investigate the impact of different configurations on the operation of the various network services, especially RRS, as many parameters from the number of probes to the pre-emption strategy can be changed.

8. Acknowledgement

This work was supported by Defence R&D Canada (DRDC).

9. References

- Andel, T.R & Yasinsac, A., On the Credibility of MANET Simulations, *IEEE Computer Magazine*, Vol. 39, No. 7, (Jul. 2006): 48-54, ISSN: 0018-9162.
- Anh, G-S. et al. (2002) Supporting Service Differentiation for Real-Time and Best Effort Traffic in Stateless Wireless Ad Hoc Networks (SWAN), *IEEE Transactions on Mobile Computing*, Vol. 1, No. 3, (Jul.-Sep. 2002): 192 - 207, ISSN: 1536-1233
- AUSCANZUKUS (1999) Multi-National Naval Task Group (MNTG) Final Report, *Naval C4 JWID Publication JWID99-R*, Washington, DC, Sep 1999.
- AUSCANZUKUS (2003) Maritime Tactical Wide Area Networking (MTWAN), *Publication ACP 200*, Washington, DC, Jul 2003
- Barsaleau D. & Tummala, M. (2004) Testing of DiffServ Performance over a U.S. Navy Satellite Communication Network, *Proceedings of MILCOM 2004*, pp. 528 - 534, ISBN: 0-7803-8847-X, Monterey, CA, Oct. - Nov 2004.
- Braden, R. et al., "Resource ReSeRVation Protocol (RSVP) -- Version 1 Functional Specification", IETF RFC 2205, Sep. 1997.

- Holliday, P. (2005). "Techniques for Efficient Network Layer Failover in Maritime Tactical Wide Area Networks (MTWAN), *Proceedings of MILCOM 2005*, pp. 52 – 55, Vol. 1, ISBN: 0-7803-9393-7, Atlantic City, NJ, Oct. 2005.
- Jorgenson, M., Reichelt, C., & Johnson, T. (2005). Operation of the Dynamic TDMA Subnet Relay System with HF Bearers, *Proceedings of MILCOM 2005*, pp. 338 – 343, Vol. 1, ISBN: 0-7803-9393-7, Atlantic City, NJ, Oct. 2005.
- Kidston, D. & Labbé, I. (2006). A Service Oriented Framework for Policy-Based Management of Maritime Mobile Networks, *Proceedings of MILCOM 2006*, pp. 1-7, ISBN: 1-4244-0617-X, Washington, D.C., Oct. 2006.
- Kidston, D. et al (2007) A Policy-Based Resource Reservation Service for Maritime Tactical Networks, *Proceedings of MMNS*, pp. 149-160, ISBN: 978-3-540-75868-6 San José, California, Oct-Nov. 2007.
- Kidston, D. & Kunz T. (2008) Towards Network Simulations Credibility: Lessons from Applying Five Key Principles, *Proceedings of MILCOM 2008*, pp. 1-6, ISBN: 978-1-4244-2676-8, San Diego, CA, Nov. 2008.
- Kurkowski, S., Camp, T., & Colagrosso, M, MANET Simulation Studies: The Incredibles, *Mobile Computing and Communication Review*, Vol. 9, No. 4, (Oct. 2005): 50-61, ISSN:1559-1662.
- Lee, S-B., Ahn, G-S., & Campbell, A. (2001) Improving UDP and TCP Performance in Mobile Ad Hoc Networks with INSIGNIA, *IEEE Communications Magazine*, Vol. 39, No. 6 (Jun. 2001): 156-165, ISSN: 0163-6804.
- Sanchez, M. & Manzoni, P. (2001) ANEJOS: A Java-based simulator for ad-hoc networks, *Future Generation Computer Systems Magazine*, Vol. 17, No. 5. (March 2001): 573 – 583, ISSN: 0167-739X.
- Sibbald, LCdr. (2004) MARPAC PacketShaper Trial Hot Wash-Up, *MARPAC HQ Presentation N60*, Washington, DC, Nov 2004.

Modelling Simulation and Optimization

Edited by Gregorio Romero Rey and Luisa Martinez Muneta

ISBN 978-953-307-048-3

Hard cover, 708 pages

Publisher InTech

Published online 01, February, 2010

Published in print edition February, 2010

Computer-Aided Design and system analysis aim to find mathematical models that allow emulating the behaviour of components and facilities. The high competitiveness in industry, the little time available for product development and the high cost in terms of time and money of producing the initial prototypes means that the computer-aided design and analysis of products are taking on major importance. On the other hand, in most areas of engineering the components of a system are interconnected and belong to different domains of physics (mechanics, electrics, hydraulics, thermal...). When developing a complete multidisciplinary system, it needs to integrate a design procedure to ensure that it will be successfully achieved. Engineering systems require an analysis of their dynamic behaviour (evolution over time or path of their different variables). The purpose of modelling and simulating dynamic systems is to generate a set of algebraic and differential equations or a mathematical model. In order to perform rapid product optimisation iterations, the models must be formulated and evaluated in the most efficient way. Automated environments contribute to this. One of the pioneers of simulation technology in medicine defines simulation as a technique, not a technology, that replaces real experiences with guided experiences reproducing important aspects of the real world in a fully interactive fashion [iii]. In the following chapters the reader will be introduced to the world of simulation in topics of current interest such as medicine, military purposes and their use in industry for diverse applications that range from the use of networks to combining thermal, chemical or electrical aspects, among others. We hope that after reading the different sections of this book we will have succeeded in bringing across what the scientific community is doing in the field of simulation and that it will be to your interest and liking. Lastly, we would like to thank all the authors for their excellent contributions in the different areas of simulation.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

David Kidston (2010). Using Modelling and Simulation to Evaluate Network Services in Maritime Networks, Modelling Simulation and Optimization, Gregorio Romero Rey and Luisa Martinez Muneta (Ed.), ISBN: 978-953-307-048-3, InTech, Available from: <http://www.intechopen.com/books/modelling-simulation-and-optimization/using-modelling-and-simulation-to-evaluate-network-services-in-maritime-networks>

INTech
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China

51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2010 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.