

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.

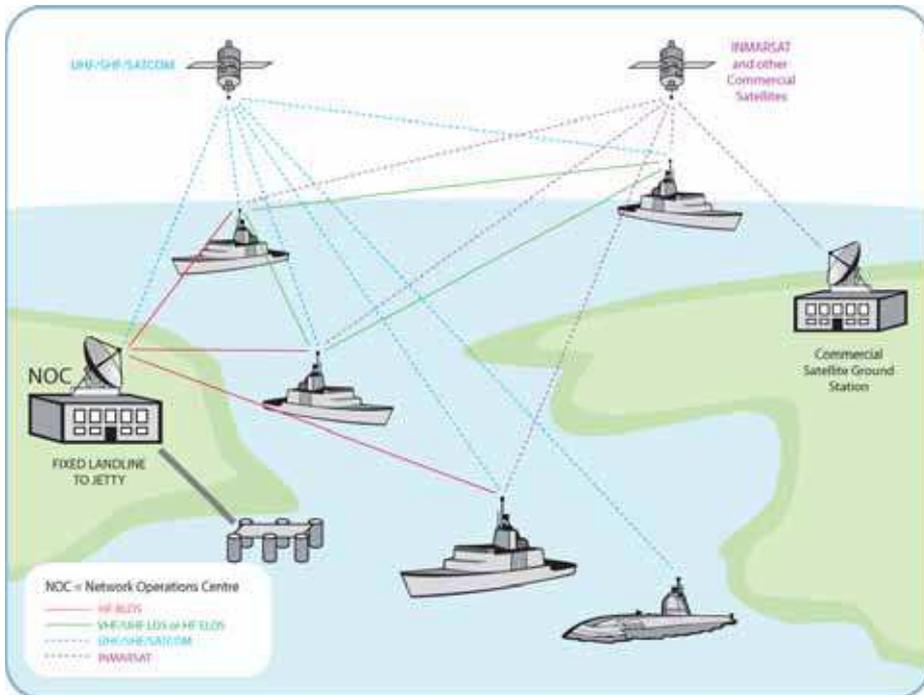


Fig. 1. Typical Maritime network

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).

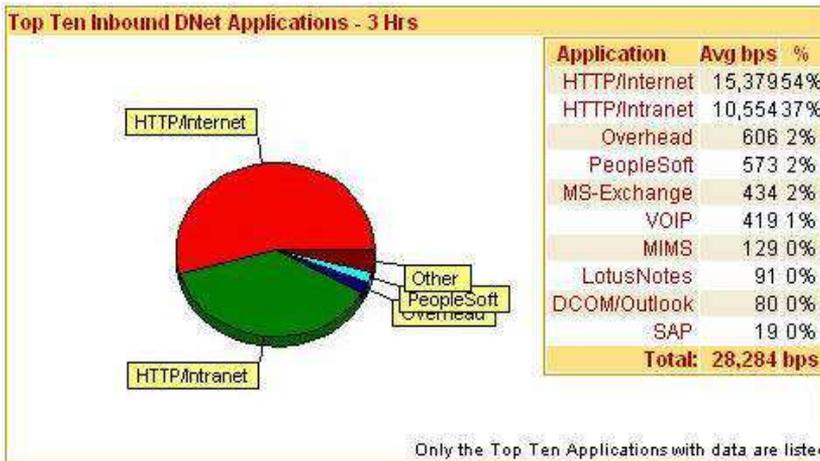


Fig. 2. Traffic Breakdown for a Naval Exercise

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.

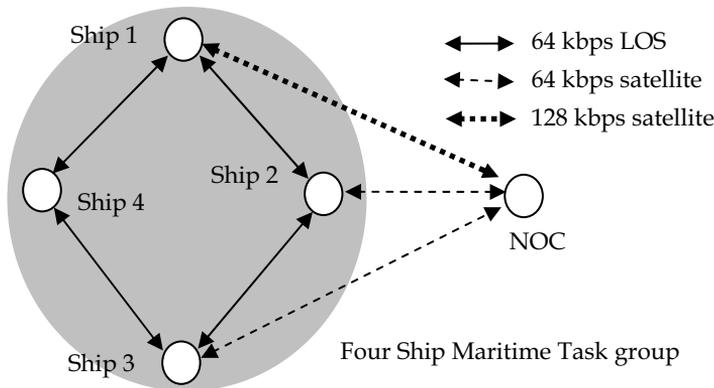


Fig. 3. Small Network

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).

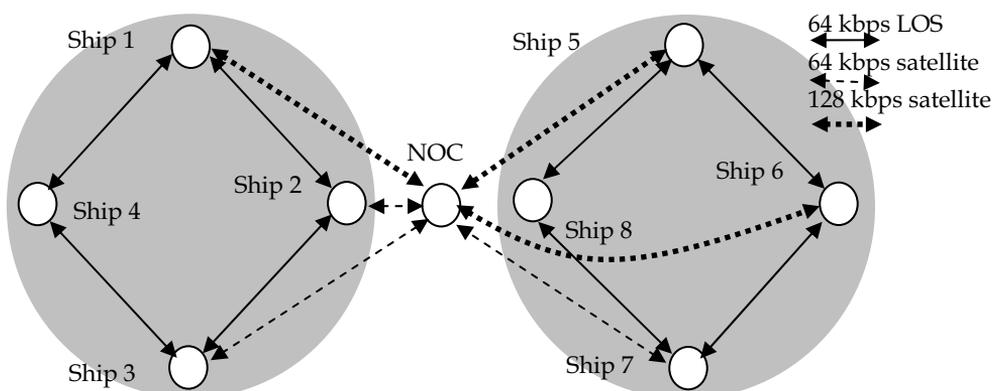


Fig. 4. Large Network

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

since the NOC is connected to mobile nodes by satellite, such links are available at all times to the nodes at which such links are operational. Since modern satellite systems can achieve MTBF rates of > 5,000 hours with MTTR of < 1.0 hour, the failure of satellite links has not been modeled.

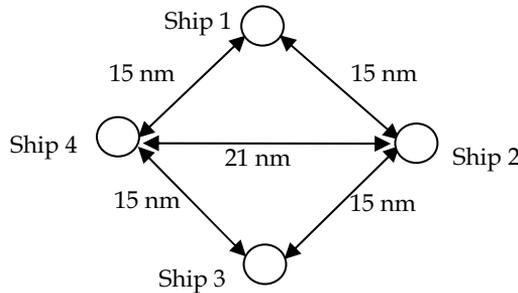


Fig. 5. Task force geometric configuration

To give an idea of the impact of this type of mobility considers ship 4 which does not have a satellite link to the NOC. The preferred gateway to NOC of ship 4 in the network was analysed. In this model, ship 4 is connected via LOS links only. The preferred gateway is ship 1 96.2% of the time, ship 3 3.7% of the time, and ship 4 is disconnected from the network 0.1 % of the time. Note that since ship 1 has a higher speed satellite link, it is therefore preferred over ship 3.

The second part of maritime mobility is inter-task group mobility, which applies only to the large network. In this model, the two task groups begin 18 Nm away from each other (at the closest point) at a random angle (from 0° to 360°). The first task group then approaches the other steadily at 30 knots (Nm/hour) on a set heading evenly distributed from this angle - 45° to +45° with 0° being directly towards the centre of the other task group. In combination with intra-task group mobility, there will be link failures and recoveries based on the 18 Nm range of the LOS links. This mobility model is outlined graphically in Figure 6.

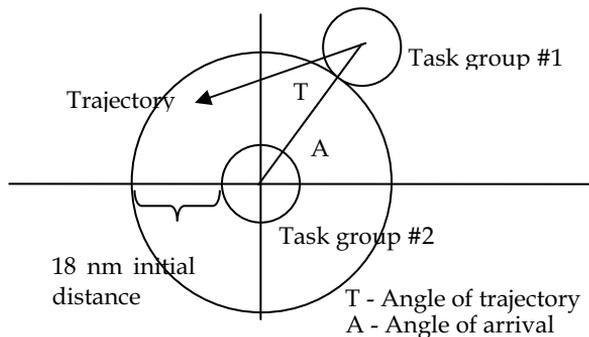


Fig. 6. Inter-task group mobility

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.

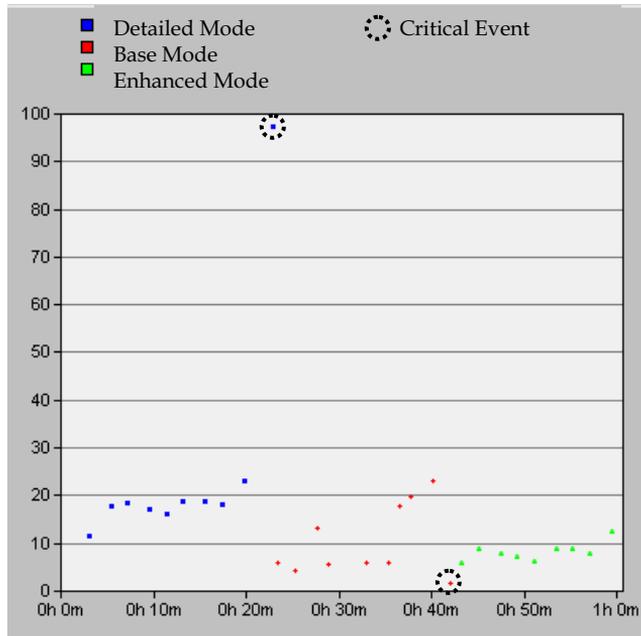


Fig. 7. Traffic monitoring adaptability example - small network - dynamic load

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

Thank You for previewing this eBook

You can read the full version of this eBook in different formats:

- HTML (Free /Available to everyone)
- PDF / TXT (Available to V.I.P. members. Free Standard members can access up to 5 PDF/TXT eBooks per month each month)
- Epub & Mobipocket (Exclusive to V.I.P. members)

To download this full book, simply select the format you desire below

