

A Policy-Based Resource Reservation Service for Maritime Tactical Networks

David Kidston¹, Isabelle Labbé¹, Francis St-Onge¹, and Thomas Kunz²

¹ Communications Research Centre
Ottawa, Ontario, Canada
² Carleton University
Ottawa, Ontario, Canada

Abstract. Naval at sea (maritime tactical) networks are characterised by a dynamic, heterogeneous, and low-bandwidth environment. There is a critical need for Traffic Engineering (TE) mechanisms to support traffic prioritisation and resource optimisation in this environment. A desirable management service in this environment is end-to-end guaranteed bandwidth for critical application flows. Solutions such as RSVP are not appropriate for the maritime environment where links are error prone and easily overloaded. This paper describes the Resource Reservation Service (RRS), a policy-enabled flow-based TE management service developed specifically for the low-bandwidth, high-error rate, and mobility of the maritime environment. This service includes several novel features including multi-path probing, bi-directional reservations, and full policy control. The value of multi-path probing is demonstrated by simulation.

1 Introduction

Traffic engineering (TE) in maritime tactical networks represents a challenge for several reasons. Maritime networks are composed of heterogeneous links that are error-prone, failure-prone, high-latency, and offer relatively low bandwidth communication capacity. For these reasons traffic should be directed over links that support their QoS requirements while making best use of the network capacity available. This task is complicated by the navy's hierarchical command structure, which requires that network management authority be partially decentralised such that a subordinate level of network control remains available to the commander of each maritime node (ship). A final complicating factor for TE in this environment is the limited availability of skilled network operators. Automation of the management of communication resources is required to minimise the skill level required from operational personnel.

Policy systems are able to provide this automation by changing network behaviour to match currently stated policy. A Policy Based Traffic Management (PBTM) system was developed to support the automation of TE management services in the maritime environment [1]. Several services were developed, including a class-based adaptive routing service (forwarding depending on what delay, error, and bandwidth the type of traffic requires) and a class-based traffic prioritisation service (a DiffServ based scheme where traffic is prioritised depending on its value to the current mission).

This paper describes another policy-enabled TE management service designed for the maritime environment, the resource reservation service (RRS). The RRS has several advantages over existing flow-based resource optimisation protocols. A distributed admission control scheme provides load balancing by probing multiple pre-computed routes at once. The use of time-outs and acknowledgements provides fault-tolerance. Bidirectional reservations are supported by making reservations in both directions at once. The RRS also coordinates its operation with the existing TE services to provide flow-appropriate routing and relative prioritisation when reservations fail.

The remainder of the paper is organised as follows. A description of the maritime environment provides an outline of TE management challenges in this area. This is followed by a description of the PBTM system which supports policy-based automation of the RRS by supporting routing and pre-emption decisions amongst other things. The four phases of the RRS operation are then described. Simulation results regarding the multi-route probing of the service follow. Finally, the paper ends with a discussion of related work, conclusions, and future work.

2 The Maritime Environment

Naval units (nodes) most commonly communicate in a combination of two modes [2]. First, they may communicate back to their strategic network using satellite communications (e.g. INMARSAT). This can be done in series to provide high bandwidth but high delay ship to ship communications. Second, ships may communicate directly with other ships via limited range radio (e.g. UHF LOS). Recently UHF/VHF relay technology has improved to the point that terrestrial radio systems may form mobile ad-hoc networks (MANETs). MANETs provide low bandwidth, low delay connectivity over a limited distance. A typical small task fleet deployment is shown in Fig. 1.




Fig. 1. Example of a Maritime Network

The maritime environment engenders several complicating factors that must be dealt with when designing a flow-based reservation service. These issues include: low bandwidth heterogeneous links; link failures engendered by node mobility; a hierarchical command structure which imposes strict but time varying traffic priority; and a lack of skilled network operators.

The RRS ensures that reservations do not use links that for reasons of limited bandwidth, delay, or error rate would not otherwise meet the QoS requirements of the requesting flow. These policies are provided by the class-based adaptive routing service (another TE management service). For instance, line of sight (UHF/VHF LOS) links pose a problem for the RRS since the media is shared and residual bandwidth cannot be reliably determined. There are currently no standards for QoS support in the Medium Access Control (MAC) layer of these media. Implementing QoS in these environments would involve probing, cross-layer communication, and/or instrumentation of the MAC with proprietary SNMP MIBs. These methods are not standardised and any attempt at direct measurement is likely to introduce significant overhead [3]. It is currently assumed that LOS links with more than 2 members within range are not suitable for reservations. When the link is dedicated and provides stable bandwidth, the residual bandwidth (the metric reserved by the RRS) can be determined by looking at the bandwidth available to be reserved and the amount currently reserved. This is the case for most satellite links where the media is not shared (i.e. a point-to-point link). Their operation can often be characterised as either available at full capacity or not available (binary). A policy-defined percentage of this nominal transmit capacity is set aside for reservations based on link type, connectivity, etc. at the discretion of a network architect. Residual bandwidth is then directly calculable by subtracting the currently reserved bandwidth from the link's reservation bandwidth pool. Note that the mechanism we use is the class-based traffic prioritisation TE service (DiffServ bandwidth sharing), so bandwidth that is unused by reservations will be shared amongst the remaining traffic classes.

The approach adopted for dealing with mobility is that reservations should be updated "periodically" and monitored "constantly." Existing reservations are probed at regular intervals. When a link on a reserved path fails, the reservation will be degraded to the policy-based default prioritisation level that was in force for that traffic prior to the reservation. This remains in effect until the reservation can be restored on a new route or the reservation is dropped and the user is informed of the error.

Reservation priority, as it is currently applied in fixed networks, reflects the needs and privileges of the particular application, workstation, or user. In the maritime environment all these factors must be considered, but in addition the current importance of the information being transported to the current mission (local importance) and the importance of the mission globally (domain importance) must be considered. The RRS uses a policy system to maintain a per flow priority based on these factors.

One final complication engendered by maritime networks is a lack of skilled operators to implement the required traffic engineering solutions. For this reason the RRS makes use of a policy system to automatically generate and distribute router configurations. The policy approach has the advantage that device configurations can be deployed quickly and in a coordinated fashion across the network. Finally the policy system provides authentication services which control who is allowed to submit reservations and at what level of priority through role-based security mechanisms.

3 The PBTM Policy System

The policy-based network management architecture upon which the PBTM is based has been previously described in [1,4] and is summarised here in Figure 2. The architecture is Web Services-based and fully distributed with all components residing on all nodes. The architecture is generic in that it is not specific to the maritime environment and can be extended to other service areas.

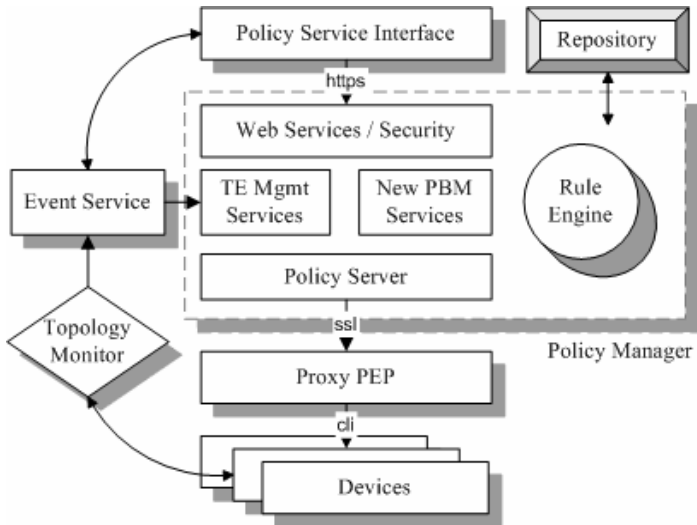


Fig. 2. Service Oriented Policy Architecture

The main components of the architecture are: the policy service interface, which accepts policy from operators while assigning roles to devices; the policy manager, which interprets high level policy and pushes low level policy out to policy-enabled resources (also known as PEPs); the proxy PEP, which takes the low level policy and configures their associated devices to conform with policy; repositories, which store the high level policies; a topology monitor, which notes changes in the network connectivity; and finally an event service, which helps to distribute events from the policy system and the underlying network (topology monitor).

As mentioned previously, the PBTM currently supports several TE services. The class-based traffic prioritisation service provides DiffServ-style traffic prioritisation for non-critical traffic. Load balancing is an important concern for resource optimisation and is enforced by the class-based adaptive routing service. Traffic prioritisation for critical flows is provided by the Resource Reservation Service (RRS). Together these management services provide an integrated solution to TE in this environment.

The PBTM system supports the RRS by providing policy supported decisions about reservation link bandwidth pool size, potential routing of the admitted flows (via the class-based adaptive routing service), prioritisation after reservation failure (via the class-based traffic prioritisation service), timeout and acknowledgement schemes, route selection, reservation restoration due to mobility or link failure, and the priority/pre-emption of reservations.

4 Main Algorithm

The Resource Reservation Service (RRS) uses distributed admission control to limit the number of flows that can use a pool of bandwidth reserved on each link in the route between source and destination. The goal of the RRS is to provide a guarantee of end-to-end QoS for a particular application flow. This sort of protection is most commonly useful for real-time applications (such as VOIP or video), but could also be used for critical data transfers (such as a specific image transfer or chat session).

The RRS consists of four phases. In the first phase, global link information is used to generate multiple routes between the source and destination of the requesting flow. The second phase of the algorithm probes the potential routes separately to determine if sufficient resources are available on all links. In the third phase, an acceptable path is selected and committed. Finally in the fourth phase the reservation is maintained until the flow terminates, 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.

Initialisation: The topology of the network is continuously updated using routing information available on the local router. OSPF, the routing algorithm of choice in maritime networks [2], regularly sends Link State Advertisements (LSAs) to distribute knowledge of the domain's connectivity information. Each router stores a complete set of the most recent LSAs in a Link State Database (LSDB). From the standard OSPF LSDB, the topology discovery module can extract: a list of current links in the domain with their associated cost metric, and node connectivity. Considering the relatively slow rate of mobility in maritime networks, the instability before OSPF can re-converge after a link failure is short compared to the operational time of the network. By using a predetermined OSPF link costs chart (an example is shown in Table 1.), the type and hence the characteristics of the links can be determined directly. Note that links with a shared medium such as UHF LOS are not suitable for reservations and are thus ignored by policy.

Table 1. Equating OSPF Cost to Link Type (an example)

OSPF Cost	Link type	Nominal BW (kbps)
750	INMARSAT	64
800	SHF SATCOM	128
1150	VHF/UHF LOS	64 (shared)
1300	UHF SATCOM	32 (shared)
1900	HF BLOS	9.6

In order to determine the bandwidth available for resource reservation on each link, the available bandwidth on each link is divided so that a percentage is assigned for reservations (which require admission control). The size of this reserved pool is based on policy and is currently 50% of the nominal bandwidth for all link types. Note that the reserved pool of bandwidth will be available for other traffic if not used by reserved traffic. This information is used for route generation and resource allocation as explained in the following subsections.

Phase One: When a user makes a request, topology information is used to generate a set of potential reservation routes dynamically. A partially disjoint routing algorithm is used to generate multiple routes in series as follows. The first route generated is the least-cost (highest bandwidth) route from source to destination (from the Dijkstra algorithm) while ignoring links that are not acceptable for the requested class of application (not enough bandwidth or counter indicated error or delay characteristics as defined in the adaptive routing policy). The second route is generated similarly, but ignores the highest cost (least bandwidth) link of the best route previously generated. The third route also uses the same algorithm, but ignores the highest cost links of the previous routes, etc. This can continue until no more routes are possible or the policy defined maximum number of routes for the reservation's priority is reached. This simple algorithm may be refined at a later date to include constraints such as error rates and latencies important in multimedia traffic. Once the routes have been generated, the generated routes are probed in parallel with each node on each route performing admission control (phase two). It is up to the destination to choose the route that will be reserved, assuming an acceptable route is available (phase three).

The main advantage of probing multiple paths is to discover the "best" path currently available. If several paths are acceptable the receiver will have a choice of selecting the reserved path such that the reservation can be made with minimal impact on the existing flows (least number of pre-empted flows). Another advantage is that probing multiple paths promotes load balancing. Where default routing forces all traffic over the "best" link, when multiple routes are considered, the route with the least loaded links can be identified for reservation. This type of selection allows the traffic to be balanced both at individual nodes and throughout the network.

Phase Two: Admission control in the proposed algorithm is similar to RSVP [5] but modified for the maritime environment. RSVP was found to be unsuitable for three reasons. First, RSVP assumes unidirectional reservations where most IP based applications are bidirectional. Second, RSVP uses the default routing to attempt reservations and does not probe multiple routes in parallel. In the low bandwidth maritime environment, the default route would be quickly overloaded and attempting alternate routes will increase the call acceptance rate. Third, although the RSVP standard has provisions for including policy control information; most implementations do not support this capability (for example, Cisco). This is required for communication with the RRS at each hop in the reservation to determine whether the flow should be admitted or not (depending on both local policy and the policy carried by the resource request). Finally it should be noted that RSVP was not designed for low bandwidth links. Its rate of signalling is high (although configurable) and it was not designed for robustness (it was mainly designed to be used over wired networks).

Instead of using RSVP for admission control, we have developed a proprietary robust signalling protocol. Admission control decisions are performed at each hop along the selected route(s). If sufficient resources exist for the desired link at the current node, the residual bandwidth of the link is noted in the probe and forwarded to the next node. Route probing is robust in that every probe message is acknowledged. Unacknowledged probes are resent after a policy configurable timeout. If a probe remains unacknowledged, the route is considered lost. Duplicate probes are ignored.

If insufficient resources are found at a node, the resources are re-checked to see if pre-empting lower priority flows would leave enough resources (pre-emption is

explained below). If sufficient resources are still not available a failure message is sent to the destination, which will in turn inform the source once all probes have arrived. If sufficient resources are available, the request is forwarded to the next node in the route taking note of the flows that would be pre-empted if this route were used.

A copy of every resource request is stored at each node in the hope that a confirmation will eventually arrive. Only at that point will the resources be committed. If the confirmation has not arrived in a policy-defined amount of time, the “pending resource request” record is purged. It is important to realise that no change is made to active reservations or the router during phase two. The purpose of route probing is simply to determine if a reservation is possible along any of the generated routes. This may lead to the case where reservations are tentatively admitted but the resources are not available when the commit packet returns, because another reservation has committed first. We argue that these false admissions are preferable to the alternative of reserving resources during the initial probing. It is more likely that bandwidth reserved during probing will be wasted because downstream nodes are not able to handle the request. In other words, routes which have been probed are unlikely to be probed and committed on the same link before the original request has a chance to return and commit its bandwidth. Similarly, only one of the multiple probed paths will be reserved so there is no point in reserving resources on links that may not be used.

A novel capability of this algorithm is that a reservation in the reverse direction (destination to source) can be made at the same time as the forward direction (source to destination). Bidirectional reservations can be especially useful when the application has critical traffic in the reverse direction that needs protecting, such as VOIP calls or FTP downloads. The bandwidth requested need not be the same in both directions. Making bidirectional reservations reduces overhead and latency while ensuring that the reservation is bidirectional (it reserves at the same nodes at the same time for both directions on two different links).

Phase Three: Once one or more reservation probes have reached the destination, phase three, route selection, is performed. Several factors may influence the choice of route. Avoiding the pre-emption of existing flows is considered. In order to balance the load of the network, the minimal residual bandwidth of each route is considered. The residual bandwidth of each link on a route was noted during the route probing phase (phase two). In heavily reserved networks the number and priority of

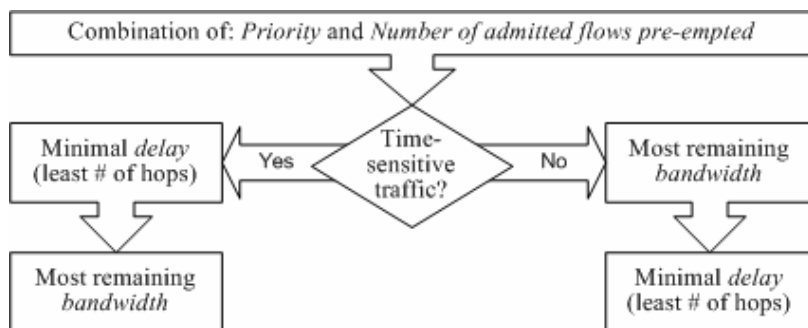


Fig. 3. Route Selection Algorithm

pre-empted flows along each route are also noted in the probes. Finally, the length of the route is also considered. The selection algorithm is presented in Fig. 3.

The selection algorithm first selects the route with the least number (and lowest priority) reservations pre-empted. If more than one route has the same number and type of pre-emptions, one of two possible selection methods are used. For real-time applications, delay is more important than raw bandwidth and therefore the flow should take the shortest route. Conversely, non real-time (data as opposed to delay focused) applications take the route with the highest minimum residual bandwidth. The rationale for this is that real-time applications are more delay sensitive than they are bandwidth sensitive. Placing them on routes that are close to saturation may be advisable if the delay is reduced. Since delay is most often a factor of hop-count, the lower hop count is favoured for these applications. On the other hand, applications with heavy bandwidth requirements are more concerned with total data transmitted, a longer delay may be tolerated and thus routes with greater residual bandwidth are preferred.

Once the destination has decided upon a route, a commit message is sent back along that route, with each RRS updating the configuration of the router so that the flow is treated in the reserved class. In order to determine if a new request should actually be committed the RRS checks which reservations are using which local links.

If insufficient bandwidth is available in the reservation bandwidth pool for the identified link, the RRS will pre-empt existing flows only if it would free enough resources to admit the committing flow. Users assign a priority to a reservation when it is made which is then modified based on the current importance of the application, data, workstation and mission according to policy. Established but lower priority calls will be pre-empted in priority order lowest to highest and then largest to smallest in terms of bandwidth.

When a reservation is pre-empted all nodes along the pre-empted reservation's route are notified to release the related resources. Depending on the policy/priority of the pre-empted or unsuccessful reservation, a request may be reattempted at the source or may simply be dropped.

Phase Four: When a commit message reaches the source the reservation enters its active phase. Each committed reservation causes the dynamic creation of an explicit MPLS tunnel to force the identified flow down the reserved path. This tunnel is torn down when the reservation ends. Policing is applied on the flow to limit its data rate to the reserved bandwidth. This is to ensure that an admitted reservation will not degrade other reserved flows. Finally, bandwidth is guaranteed at each node end to end by defining a unique traffic class.

Reservations can be terminated by a number of events including termination by the user, end of the reserved period, pre-emption by a higher priority flow, or failure of a link on the reserved route. During the maintenance phase, keep-alive messages are sent along the reserved path at a policy-defined interval. Each RRS instance on the path must receive one of these messages within another policy-defined interval or the reservation is considered to have terminated. This "natural" termination causes the RRS to simply remove the reservation from its list, freeing the associated bandwidth and reversing any router configuration that has been made.

If a link fails or is degraded, due to mobility for instance, existing reservations must react to the changed topology. Fault management is achieved by eventually terminating reservations that use the failed link and potentially attempting to re-establish

those reservations at the source (based on policy/priority). When a link is flagged by the routing protocol as no longer in the topology, the RRS will wait for a policy configurable amount of time before reacting to allow the link some time to recover. All reservations that are currently using a failed link will then be released. If the link is simply degraded (defined by policy), after the waiting period the service will recalculate whether admitted flows can still receive the resources they reserved on a degraded (reduced bandwidth) link. If not, pre-emption of the lowest priority flows continues until there is sufficient bandwidth available for the remaining reservations.

5 Simulation Results

The commercial tool OPNET [6] has been used to simulate the Resource Reservation Service (RRS). In an initial study of performance, the multi-probing feature of the RRS was measured including overhead introduced by the RRS, call acceptance, and pre-emption rate in a small (5 node) and larger (9 node) simulated maritime network. The (simulated) small network consists of four routers and four workstations and the larger network of eight routers and eight workstations. The routers use the base Cisco 3640 model with a modified OSPF stack that forwards the LSA graph to the appropriate RRS process in its associated LAN workstation. The workstations use the base Intel advanced model with the addition of a RRS process and packet generation process (used to initiate new requests locally). The routers are connected to their associated workstation with 10BaseT Ethernet and to each other by point to point links with bandwidth between 32 and 128 kbps as shown in Fig. 1. (for the small network). The larger network is composed of an additional 4 ships connected to the NOC. These network sizes and topologies were chosen based on current maritime deployments [2].

In order to remove the effects of a particular seed value, five runs at each load level are performed with different seeds. Statistics are averaged over those seeds. Simulation runs of 10,000 seconds have been used with 3 priority levels, 1 or 3 parallel probes, and with a reservation inter-arrival time for the network exponentially distributed and centered on 15s for low load and 7.5s for high load. These values were chosen to emulate a saturated network (low load) and an overloaded network (high load). Reservation sources and destinations are uniformly distributed between nodes. All reservations are for 8kbps with 50% of link bandwidth available for reservation. Reservations lifetimes are exponentially distributed with a mean of 270 seconds. In order to focus on the multi-routing aspect of the RRS, mobility has been disabled for this particular set of simulations. The results of the simulations are shown in Table 2.

The per-request (network wide) reservation overhead in bytes/call is high because it currently carries a complete reservation policy (encoded in XML) which is on average 3015 bytes of data. Compression of the XML document would improve both overhead and reservation setup time, which currently takes 2.5-3.5 seconds on average because of the low bandwidth environment. Reservation overhead includes each IP packet involved in a single request summed over each link on which it is sent. The average overhead of successful requests is shown. Note the difference in overhead between the low and high load of requests. The lower total overhead in high load networks is likely because when links become congested with requests, longer routes are rejected immediately at the source and thus the reservation request does not have to travel over as many links lowering the per-request overhead.

Table 2. Multi-Probing Simulation Results

Network Size	Load	Number of Routes Probed	Reservation Overhead (bytes/call)	Call Acceptance Rate (%)	Call Pre-emption Rate (%)
Small	Low load	3	13452	89.0	7.4
		1	3565	87.0	10.8
	High load	3	11114	69.6	19.4
		1	3544	71.7	20.4
Large	Low load	3	19028	87.0	7.4
		1	5836	83.2	7.6
	High load	3	15887	68.4	14.6
		1	5309	64.1	13.3

The call acceptance rate at low load (saturated network) is acceptable for the maritime environment, but not ideal. Even at high load (congested network) more calls are accepted than rejected. Further work will investigate the value of priority in pre-emption to give the most critical flows an event higher acceptance rate (where 99.9% would be more appropriate). As expected, the call acceptance rate is higher and call pre-emption rate lower for lower request loads. Also as can be expected, the pre-emption rate is higher for higher loads. The difference in these values based on the number of probes is almost negligible in the small network where there is a lack of alternate routes. However, in the larger network a definite advantage can be seen in improved call acceptance rates at the cost of higher overhead. Since RSVP is similar to RRS with a single probe, it is expected that this will remain the case. An investigation to directly compare the RRS and RSVP is currently being pursued.

6 Related Work

There has been very little research done to date investigating the TE requirements of maritime networks. A notable exception has been recent work in applying static Diff-Serv QoS to maritime networks [7], which showed that throughput and delay guarantees were hard to achieve in this environment. However, queuing and dropping mechanisms, if properly tuned, could provide limited service differentiation. The static nature of the DiffServ marking does not respect the dynamic nature of the maritime environment, where the importance attached to different information flows vary with time. It is however straightforward to implement and maintain. The RRS uses a policy system to accommodate changes in priority by automatically modifying the reservation priority of flows to meet their current mission value (as defined by policy).

RSVP-TE [8], a standard from the IETF, defines a traffic engineering enhanced version of the reservation protocol RSVP. RSVP-TE provides a mechanism by which MPLS label switched tunnels can be configured along a predetermined (explicit) route with or without a resource reservation being made at the same time. RSVP is used as a signalling protocol that can create and reroute label-switched tunnels. Rerouting may be required to bypass networks failures, congestion, and/or network bottlenecks.

If reservations are made, they can be pre-empted by higher priority reservations. All this is similar to the RRS. However the method for choosing a label-switched path is not specified in RSVP-TE. Also, reservations are unidirectional and are made in a single pass as a reservation message passes from node to node along the selected path from receiver to sender. The RRS allows bidirectional reservation to be made to ensure the delay/error characteristics of the label-switched path are symmetrical (an important characteristic for some types of flows). Finally, our work investigates a mechanism by which multiple potential routes are considered before selecting the most policy acceptable in order to increase the chance of call being accepted and to balance reservations over the network. RSVP-TE only considers the default route.

In terms of the policy system, hierarchical policy-based network management systems have been shown to be capable of providing user-configurable monitoring and automated configuration of MANETs [9]. The primary management services in their work include monitoring, data aggregation, and reporting. Our work also investigates the use of policy in network management in a similar environment but it focuses on TE management services.

7 Conclusions and Future Work

This paper describes the Resource Reservation Service (RRS), a policy-enabled flow-based bandwidth reservation service designed to support end-to-end QoS in maritime networks. Existing reservation schemes such as RSVP are not appropriate since links are error prone and easily overloaded. This service in addition to traffic prioritisation and adaptive routing management services are designed to provide Traffic Engineering (TE) mechanisms to support traffic prioritisation and resource optimisation in this heterogeneous, mobile, and low bandwidth environment.

The RRS provides several novel features to improve resource reservations in this environment. Since the RRS makes use of topology information available at every edge router to determine link types and connectivity, no additional overhead is required to generate routes from source to destination. Policy control ensures that multiple potential routes are generated that traverse links with sufficient raw bandwidth and have delay and error characteristics acceptable for the traffic type (according to operational policy). These routes are probed in parallel to increase the chance that a route will be found. When multiple acceptable routes are found, the route reserved is chosen to make the least impact on existing traffic, causing reservations to be balanced across the network. The use of acknowledgements, timers, and a retransmission scheme are used to mitigate the dynamic and error prone environment. For the same reason, the ability to make reservations for traffic in both directions at the same time is an advantage both to ensure bandwidth is simultaneously available and to reduce the time/bandwidth overhead compared to sequential reservations.

Simulations of the multi-probing aspect of the protocol show that, as would be expected, the call acceptance rate is higher and the call pre-emption rate is lower for lower request loads. Also, a lower overhead was observed with higher load, likely because as links become congested with requests, longer routes are rejected and reservations do not travel over as many links. The benefit of probing multiple paths becomes apparent in larger networks where there is a higher call admission rate.

While this service provides the basic mechanisms required for an efficient and robust TE service for critical flows in maritime networks, there are still many avenues for future work. Currently resource requests are unicast (one source and one destination). A multicast resource request may also be possible if RSVP-type mechanisms were to be used to merge multiple reservations, a useful capability for video conferencing and other broadcast communications. Another potential enhancement would be the use of MPLS route protection mechanisms to support the reservation of alternate parallel paths for very high-priority flows. The alternate routes on “warm standby” would be reserved in advance in anticipation of link outages with an immediate switchover when the primary reserved route fails. Finally, a more comprehensive evaluation of the RRS is planned by comparing the RRS with alternative resource reservation mechanisms such as RSVP [5] (designed for fixed networks) and INSIGNIA [10] (designed for MANETs).

Acknowledgements

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

References

1. Kidston, D., Labbé, I.: A Service Oriented Framework for Policy-Based Management of Maritime Mobile Networks. In: MILCOM 2006, Washington, D.C., USA (October 2006)
2. Maritime Tactical Wide Area Networking (MTWAN), ACP 200 Project, AUSCANZUKUS Unclassified (available through author) (July 2003)
3. Jorgenson, M., Reichelt, C., Johnson, T.: Operation of the Dynamic TDMA Subnet Relay System with HF Bearers. In: MILCOM 2005, Atlantic City, NJ, USA (October 2005)
4. Labbé, I., St-Onge, F., Kidston, D., Roy, J.-F.: A Policy System for Traffic Management in Maritime Tactical Networks, DRDC Technical Report TR-2007-005 (January 2007)
5. Braden, R., et al.: Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification, IETF RFC 2205 (September 1997)
6. OPNET: web site, (last accessed July 11, 2006), <http://www.opnet.com/>
7. Barsaleau, D., Tummala, M.: Testing of DiffServ Performance over a U.S. Navy Satellite Communication Network. In: MILCOM 2004, Monterey, CA (October-November 2004)
8. Awduche, D., et al.: RSVP-TE: Extensions to RSVP for LSP Tunnels, IETF RFC 3209 (December 2001)
9. Chadha, R., et al.: Policy Based Mobile Ad Hoc Network Management. In: POLICY 2004 (June 2004)
10. Lee, S-B., Ahn, G-S., Campbell, A.: Improving UDP and TCP Performance in Mobile Ad Hoc Networks with INSIGNIA. IEEE Communications Magazine, pp. 156–165 (June 2001)