Development of routing protocols for the Solar-powered Autonomous Underwater Vehicle (SAUV) platform

Radim Bartoš[1], Steven G. Chappell[2], Rick J. Komerska[2], Matthew M. Haag[3], Sai S. Mupparapu[2], Emmanuel Agu[4], and Ian Katz[1,2]

[1]Department of Computer Science, University of New Hampshire, Durham, NH 03824, USA
[2]Autonomous Undersea Systems Institute, Lee, NH 03824, USA
[3]Technology Systems, Inc., Brunswick, ME 04011, USA
[4]Department of Computer Science, Worcester Polytechnic Institute, Worcester, MA 01609,USA

Abstract

The Solar-powered Autonomous Underwater Vehicle (SAUV) platform has been developed over the past several years as a means to achieve long-endurance cooperative missions, and the design of mobile routing protocols has become one of the main areas of focus during this process. This paper reports on the lessons learned from the design, simulation-based evaluation, and field testing of two protocols. Two areas for possible improvements are proposed and evaluated using simulations. Finally, we present suggestions drawn from field experience for improving the design, simulation, and evaluation of protocols for use in a harsh underwater environment.

Keywords: Ad Hoc Routing; Underwater Networks; Autonomous Underwater Vehicles; Unmanned Underwater Vehicles.

1 Introduction

Ad hoc routing protocols are of particular interest in the domain of Autonomous Underwater Vehicles (AUVs), as they enable communication among nodes that are not within range of a direct transmission; this enhances the ability of AUVs to work cooperatively underwater. There are existing ad hoc routing protocols in the domain of RF wireless communications, and their performance has been extensively evaluated. Many of the current approaches to underwater acoustic networking are based on the design of these RF routing protocols, including their methods of coordination: the exchange of control messages.

Based on our field experience, however, we find that the underwater environment is significantly more challenging than has been assumed in much of the previous simulation of these protocols; even short-term stability of a communication channel cannot be taken for granted. This, combined with high signal propagation latency and low data rates, severely limits the applicability of such protocols to fleets of AUVs – the exchange of control messages is both expensive and unreliable. As a result, we believe that *robustness* should be the main focus of research on underwater routing protocols, instead of performance or efficiency improvements; proposed routing protocols must be able to operate in environment where significant packet loss and asymmetric links are a common occurrence. In this paper we present the steps taken from the initial design of a routing protocol through in-water testing to our current work on the second-generation routing schemes.

Over the past few years, the Autonomous Undersea Systems Institute (AUSI), together with its partner organizations, has been developing the Solar-powered Autonomous Underwater Vehicle (SAUV) platform [1]. To support above-water communication, the SAUV contains a FreeWave RF modem and an Iridium satellite modem. For its underwater communication, the SAUV contains a Benthos Acoustic Telemetry Modem (ATM-885PCB). This paper discusses results of communication over the acoustic modems only, which were operated at the default rate of 800 bps.

SAUVs were designed for deployment in fleets of vehicles that cooperate on a common mission. In such scenarios, the ability to support reliable multi-hop communication is critical. Two ad hoc routing protocols have been designed to address this need: Autonomous Undersea Systems Network (AUSNET) [2, 3] and Controlled Flooding for Small Networks (COFSNET) [4]. Both protocols are tailored versions of existing ad hoc routing protocols. The AUSNET effort builds upon the emerging ad-hoc (self-forming, self-maintaining) network protocol Dynamic Source Routing (DSR) [5]. COFSNET uses the concept of a controlled flood technique to facilitate multi-hop communication. It is the bare minimum protocol capable of supporting multi-hop communication in a dynamic network. This paper presents a summary of the experience obtained in the development, implementation, evaluation, and field testing of these routing protocols.

The rest of this paper is organized as follows: Section 2 relates previous and relevant work in the area. Section 3 introduces the details of the AUSNET and COFSNET protocols. Section 4 describes the various simulators employed in our work. Section 5 presents the initial set of simulations on those protocols and discusses field tests using them with multiple SAUVs. Section 6 presents lessons learned from those simulations and deployment trials. Section 7 introduces Packet Deprecation and COFSNET+: COFSNET with improvements, designed in the face of those lessons presented in Section 6. Section 8 presents results obtained while running Packet Deprecation and COFSNET+ in simulation. We present our conclusions in Section 9 and relate areas for future work in Section 10.

2 Previous Work

Underwater communication is an area that has received significant attention in recent years [6, 7, 8, 9]. Underwater networking problems are studied at many different levels starting with the design of acoustic communication links (e.g., [10]). Large propagation latencies of the acoustic signal make the design of media access control (MAC) methods very difficult [11, 12, 13, 14, 15, 16]. Many of the standard approaches to MAC have been tried, but none yet have emerged as a clear winner.

Multihop routing in underwater networks has been studied by several groups. Xie and Gibson [17] describe a mechanism through which a centralized server (gateway) manages an adaptive and self-configuring acoustic network. While centralization allows for a greater degree of optimality and manageability, it is less suitable for fast-changing AUV networks. Pompili, Melodia, and Akyildiz [18] proposed two routing algorithms, one for delay-sensitive and one for delay-insensitive applications. The work focuses on static networks and on optimization of routes based on the properties of the underlying acoustic links. Another work by this group [19] points out the need for robust self-configuring communication strategies in networks of AUVs. Carlson et al. [20] proposed a location-aware source routing (LASR) protocol and compared the performance with DSR and limited flooding. LASR assumes symmetric links, low-drift clock allowing one-way ranging, and modems that report frame error rate estimates. Under these assumptions, their simulator study shows that LASR outperforms both flooding and DSR. From the perspective of this paper, it is important to note that the simulation study indicates that the benefits of both LASR and DSR diminish at the points of the simulation with the highest rate of topology change and the smallest degree of network connectivity. Harris and Zorzi study the impact of bandwidth-distance relationship of an acoustic channel on energy-efficient route selection [21].

At the practical level, there are three significant efforts to deploy underwater networks that rely on multihop transmissions: the Seaweb project [22] (NPS and SPAWAR), the PLUSNet project [23] (WHOI), and DARPA's Collaborative Networked Autonomous Vehicles (CNAV) program [24].

3 The Protocols

Two ad hoc routing protocols have been implemented on the SAUV platform: one that augments an existing mobile ad hoc routing protocol, and one that strives for a simplistic yet functional design.

3.1 A DSR-based protocol: AUSNET

The Autonomous Undersea Systems Network (AUSNET) protocol was developed under a National Science Foundation (NSF) funded STTR Phase 2B program addressing network protocols for underwater communications. As such, AUSNET's goal is to enable expanded networking services specially tailored to the acoustic environment and AUV operational scenarios.

AUSNET builds upon Dynamic Source Routing (DSR), an existing ad-hoc routing protocol designed for mobile networks. DSR allows the network to be self-forming and self-maintaining through two mechanisms: Route Discovery and Route Maintenance. Nodes are able to learn network routes on demand as well as listen to the route discoveries of their neighbors. This allows the network to be more flexible under mobility without the need to flood the network with routing table updates.

When the route to a destination node is unknown, or sent packets are not acknowledged after a certain number of attempts, DSR will begin Route Maintenance. Outgoing packets to that node are enqueued until the Route Discovery operation has completed. The default AUSNET settings are to retransmit a packet at most 5 times (in 30-second intervals) before assuming that the route hop being attempted is no longer functional. These values define how DSR diagnoses broken routes and reacts to discover new routes.

To make DSR more suitable for underwater communication, AUSNET employs Prediction Based Routing (PBR) [25]. PBR assumes that the movement of underwater vehicles is not random; in the cases where it is not known beforehand, it can be predicted. The PBR component of AUSNET tries to make use of available movement information by making an estimation of the current network topology, based on each known vehicle trajectory and the principle of dead reckoning. Once the network topology is determined, a spanning tree algorithm provides the shortest path to the destination, thus reducing the overhead of route discovery in DSR. If the prediction is unsuccessful, AUSNET falls back on DSR's Route Discovery mechanism.

3.2 A flooding-based protocol: COFSNET

The COntrolled Flooding for Small NETworks (COFSNET) protocol [4] is a straightforward implementation of the concept of controlled flooding: a simple, coordination-free networking protocol.

Flooding is a mechanism by which a source node that needs to send a packet to a destination node simply broadcasts it. Along with a destination address, the packet contains a Time To Live (TTL) value. Any intermediate node that receives this packet will decrement the TTL and then rebroadcast it, until the TTL reaches zero. Controlled flooding (sometimes referred to as limited flooding) uses packet sequence numbers to ensure that a packet is rebroadcasted no more than once by any of the nodes in the network.

Although flooding is deservedly considered inefficient in most above-water wireless applications, that is not necessarily the case for underwater acoustic applications since the number of nodes remains relatively small and the cost of coordination could be prohibitive (such as the topology discovery process in more sophisticated routing protocols).

Unlike AUSNET, COFSNET does not determine when routes to nodes do not exist. As such, it cannot enqueue outgoing packets when their destination is unreachable.

4 Simulator Environments

Three different simulation environments were developed together with the protocols, each addressing different goals that arose during the effort. Significant attention was focused on protocol implementation fidelity: two of the simulators use the actual protocol code libraries used in the SAUV.

4.1 TSI Mobile Ad Hoc Undersea Network Simulator

The *TSI Mobile Ad Hoc Undersea Network Simulator* models mobile AUVs traversing set waypoints within a two-dimensional area, while simultaneously modeling the message movement through a transmission medium (defined by propagation speed and range). Packets can be sent individually or at regular intervals, to either random or specific hosts.

This simulator uses the SAUV's library functions, and is designed to function at approximately 30 frames per second. The simulation can be slowed down significantly in order to allow closer observation of the interactions between individual packets, or accelerated to allow for the generation of large sets of statistical data.

4.2 Discrete Event Simulator

The event-driven simulator was designed specifically to efficiently gather statistical data from a large number of runs. Unlike the TSI simulator, it does not use the SAUV's library due to the differences between the way control flows in the vehicle and in the simulator. Testing, however, has been performed to ensure the fidelity of the data generated with this simulator.

4.3 CADCON

The *Cooperative AUV Development Concept (CADCON)* is a real-time environment simulator with selected environmental interactions, which utilizes virtual machines running complete SAUV systems [26]. This setup provides maximum fidelity at the expense of simulation time, as CADCON is locked to a real time clock. Although we did not evaluate network protocol performance using this simulator, it was invaluable in staging field tests.

5 Performance Evaluation

5.1 Simulation

The goals of the simulation studies were to test the functionality of the AUSNET and COFSNET implementations before costly in-water testing, and to determine the relative merits of each. The number of nodes used in the simulations corresponds to the expected number of cooperating AUVs in typical deployments. Two quantities were measured: the total number of bytes transmitted by the nodes (relating to the total amount of energy consumed by the vehicles), and the average message latency.

A number of scenarios were tried in the TSI simulator. In this section we outline the results for the Bowtie scenarios shown in Figure 1, run using the TSI simulator. In these trials, all but the destination node were placed along a straight line equally-spaced so that just the neighboring nodes are in the communication range. The destination node C starts within communication range of the source node A and then moves in a way that forces it to use the remaining nodes as relays.

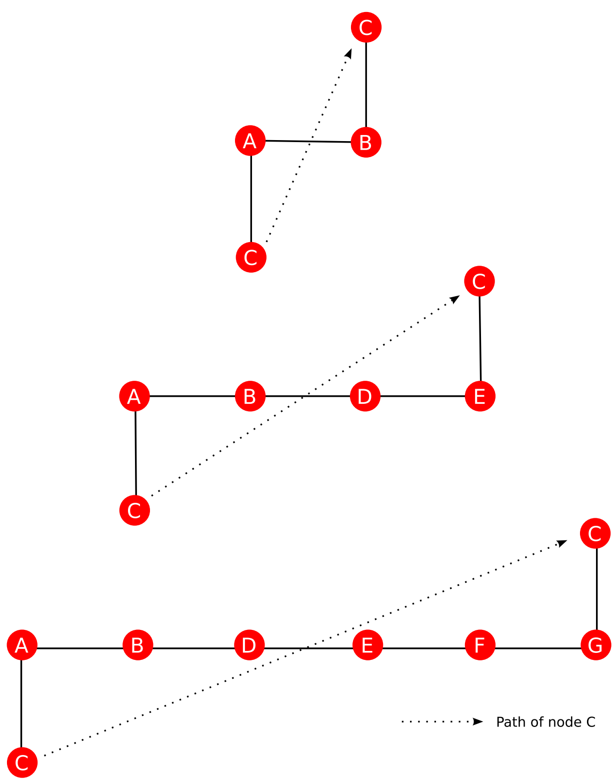

Figure 1: Bowtie scenarios for 3, 5, and 7 nodes: The path travelled by node C causes changes in routes, but no breaks in connectivity.

Figure 2 shows the overhead – the difference between the total number of bytes transmitted in the network and the actual data bytes delivered – for both protocols. COFSNET shows lesser overhead in the scenarios with 3 and 7 nodes and slightly greater in the case of 5 nodes. The COFSNET overhead is mainly due to the flooding nature of the protocol, while the AUSNET overhead is due to the route discovery packets and its larger header size.

Figure 3 shows the average latency – the difference between source node transmission time and destination receive time. Again, COFSNET performed better with relatively lower latency in all the cases. This is due to the route discovery process in AUSNET where data packets must be queued until a route is found.

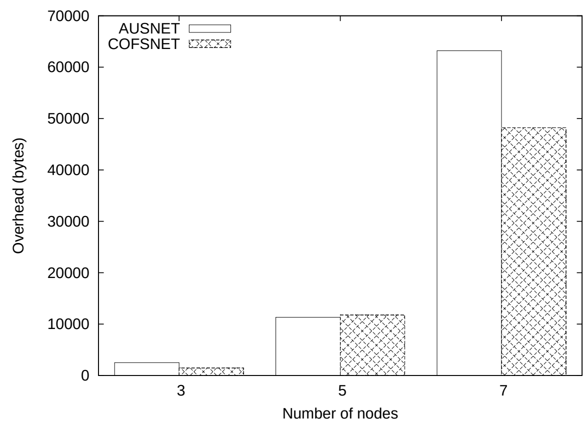

Figure 2: Protocol overhead comparison.

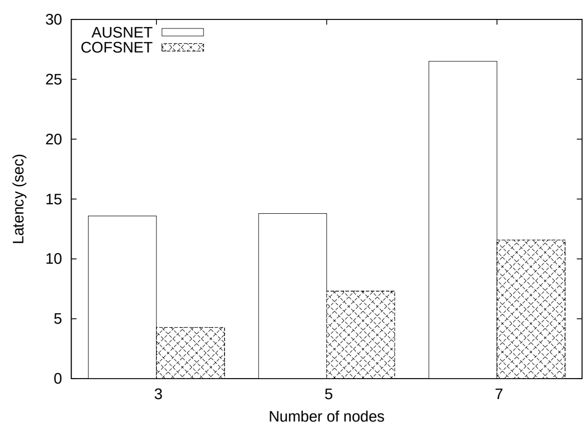

Figure 3: Message latency comparison.

Armed with these results showing acceptable performance for both AUSNET and COFSNET, we moved them to in-water trials.

5.2 Field Testing

In-water testing (see Figure 4) of the AUSNET and COFSNET protocols was performed on two occasions: in Lake George near Bolton Landing, NY in 2004 [27], and at the Office of Naval Research (ONR) sponsored AUV Fest 2005 held in Keyport, WA [28]. COFSNET has also been utilized in more recent field tests [29], although specific network performance data was not collected.

Figure 4: Preparation for in-water networking tests.

5.2.1 Lake George Test Summary

The objectives of this series of tests were to demonstrate the functionality of AUSNET and COFSNET in a real-world environment, including basic connectivity and network reconfiguration with 3 nodes in water. The nodes included two communication buoys and a Benthos "deck box" mounted on a chase boat.

The tests were conducted in the vicinity of the Dollar Islands (Figure 5) so that they could be used to block acoustic signals, a convenient alternative to moving the nodes out of range of one another. This method of breaking connectivity was verified through a separate test using modem ranging.

The communication buoys were positioned with buoy 1 between West Dollar Island and the mainland and buoy 2 in open water, so that they could not communicate directly. The chase boat was situated in a location north of the islands so that it could become the relay node between the buoys.

Six connectivity experiments were conducted by sending short ASCII messages with sources and destinations as described in Table 1. In order to test reconfiguration, buoy 2 was moved northward until it formed a line of sight connection to buoy 1, and routes could be established. Buoy 1 was then moved southward to break its connection, and the chase boat was moved northward to act as the relay node. Finally, buoy 1 was again moved to a position with a line of sight connection to buoy 2.

Table 1: Connectivity Tests: message origin and destination

Experiment	Origin	Destination
1	Boat	All nodes
2	Boat	Buoy 1
3	Boat	Buoy 2
4	Buoy 1	Buoy 2
5	Buoy 1	All nodes
6	Buoy 1	Boat

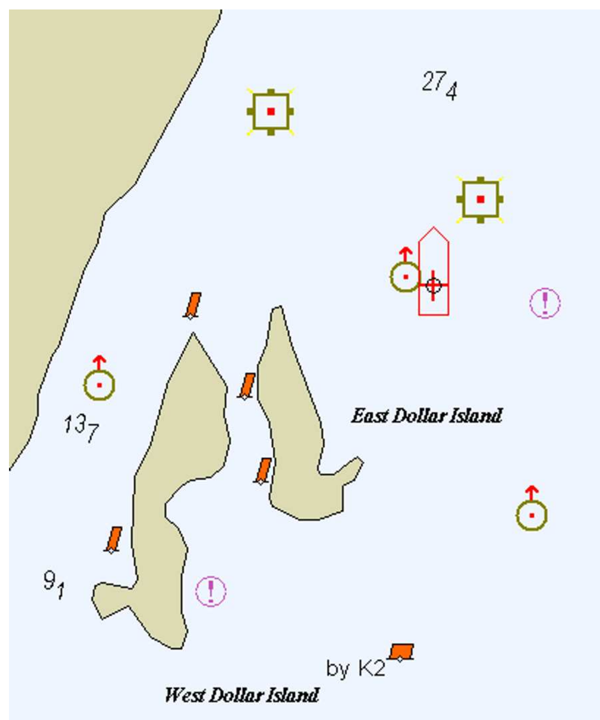

Figure 5: Lake George test setup: circles represent location of communication buoys, square icons represent locations used to test communications.

For each of these experiments, the AUSNET and COFSNET protocols performed as predicted [27]. Both protocols demonstrated connectivity and the ability to reconfigure themselves; they correctly added the relay node to the network path when direct communication was interrupted, and removed it when direct connectivity was restored. However, the added discovery time required by AUSNET when changing the network topology was apparent when compared to COFSNET, which reacted to changes immediately.

5.2.2 AUV Fest 2005 Test Summary

The objectives of this series of tests were to demonstrate the operation of AUSNET in a more heterogeneous environment of nodes. Besides the AUSI SAUV, two other types of nodes were present.

1. The Mid-Sized Autonomous Research Vehicle (MARV), maintained and operated by the Naval Undersea Warfare Center (NUWC), Newport, RI.
2. A Gateway Buoy, a simple layer 2 bridge between an RF and acoustic modem that allows operators on land to appear as a node in the underwater network.

The tests were conducted within a rectangular working area of 1.3×0.4 miles. The nodes were operated in the Cooperative Survey scenario shown in Figure 6, with two SAUVs designated as nodes A and B and the faster MARV as node C (the Gateway Buoy, connected to nodes A and B, is not shown here).

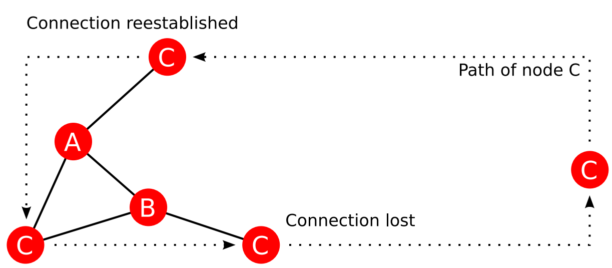

Figure 6: Cooperative Survey scenario: The path travelled by node C brings it in and out of communication range with nodes A and B.

The AUSNET protocol was initially used to provide a network layer for transmission of vehicle commands and status between the operator and vehicles. However, rapidly changing acoustic conditions in the water and unidirectional communication led to quickly shifting routes – causing serious operational issues.

The instability of the physical medium ultimately led to long interruptions in AUSNET communication. This difficulty was alleviated to some extent via changes to the transmission settings of the Benthos Acoustic Telemetry Modems, such as lowering the baud rate and transmission power. However, the communications were further hampered by excessive collisions; a lack of access control to the physical medium resulted in significant signal interference and packet loss. The results of using AUSNET to conduct these tests were deemed less than acceptable, and all nodes were updated to use the COFSNET protocol instead.

COFSNET performed surprisingly well during the testing, successfully filtering and forwarding messages as appropriate [28]. It was also noted that the communications window to node C was much wider with COFSNET than with AUSNET, due to COFSNET's faster reaction to topology changes – an essential feature in this scenario.

6 Lessons Learned from the First Development Cycle

The effects of random packet loss, asymmetric links, and node disconnections on network protocol performance were severely underestimated in our design and simulations. Not only did the overhead for the "smart" AUSNET protocol increase beyond that of the "inefficient" COFSNET protocol, the longer reconvergence times for AUSNET reduced the window of opportunity to communicate with a passing node. To be considered practical, a protocol must be able to gracefully handle disconnection and quickly recognize reconnection.

Controlled flooding is admittedly an inefficient routing protocol, but in a real environment with harsh conditions, it has proven itself as a viable option over what would ordinarily be considered a superior protocol design. Furthermore, typical AUV deployments consist of a relatively small number of nodes; scalability issues of a flooding-based protocol are of lesser significance.

6.1 Packet Queuing Problem

One of the main contributors to AUSNET's problems in the AUV Fest 2005 testing was determined to be the behavior that results when a desired destination node is unreachable. According to the initial AUSNET design, if a route is unreachable (or nonexistent) when packet transmission is attempted, the packet is placed in a queue awaiting a proper route. This seemingly reasonable behavior was found to introduce severe consequences in the network.

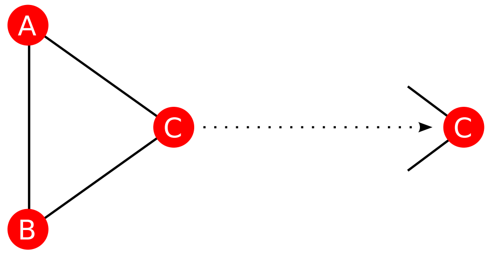

Figure 7: Node C periodically sends status packets, moves away and becomes disconnected.

To illustrate, consider the scenario in Figure 7: the AUV represented as node C is within range of nodes A and B and communicates its status once per minute. It then travels away from the other two nodes, until it becomes disconnected. During 20 minutes of disconnection, node C will accumulate 20 packets in its queue. When node C then rejoins the network, it will then consume the acoustic channel until all packets in its queue have been sent. Such channel "hogging" raises the probability that packet collisions with other nodes will occur.

While this status information is important, packets that are queued later contain updated status information and thus deprecate earlier packets. In order to save valuable bandwidth, it would be beneficial to dequeue deprecated packets and send only the most recent status packets.

6.2 Scalability

Eventually, the cost of AUV technology should fall low enough that fleets of AUVs can become large enough to face problems of network scalability. As such, COFSNET is at best a temporary solution in the long term and should give way to more efficient protocols that do not sacrifice the robustness of flooding. Due to the simplicity of the COFSNET protocol, a medium-term solution will be to extend its current functionality to address these concerns.

7 Improving Network Behavior

After completing the first development cycle, the goals for the next phase are to carry forward the successes of the protocols in their current forms and apply the insights of in-water testing to address their shortcomings. In this section, we present improvements to both protocols. These augmentations will be evaluated in Section 8.

7.1 Improving performance under disconnection: AUSNET Packet Deprecation

Outgoing packets generally fall into two distinct categories: those that are replaceable over time, and those that are not. Packets that are replaceable – such as network or node status messages – should be dropped from an output queue when a more recent version is added.

A method of preventing the accumulation of stale packets in AUSNET's outgoing queue is addressed by the proposed [25] *Replace* and *Hold* modes of packet expiration, a very simple solution. Although there are other expiration models which may be supported, such as [30] which gives the application fine-grained control over this behavior, Replace and Hold is a reasonable first step since it is the simplest.

Any packet replacement algorithm introduces additional processing. However, given the low rate of packet arrivals (constrained by the low bit rates of the acoustic modems), additional processing does not add a significant burden.

This proposed extension addresses only some of AUSNET's behavioral issues in a harsh acoustic environment. A broader solution to the problems addressed in Section 6 is the subject of future work.

7.2 Improving scalability: COFSNET+

If we consider the total cost of delivering a packet to be the number of times a copy of the packet is transmitted, then the cost of limited flooding in a contiguous network becomes the number of nodes (each node will retransmit the packet exactly once). In the original COFSNET, this cost does not depend on the relative distance between the source and the destination, nor on the topology of the network. COFSNET+ aims to improve this aspect of COFSNET by reducing the number of overall transmissions required to deliver a packet – without sacrificing its benefits [31].

When no information about the network topology is available, COFSNET+ employs limited flooding. As the data packets are delivered, network topology information is collected and used to limit the region in which rebroadcasts will take place. If errors in delivery are detected, such as those caused by a change in network topology or acoustic conditions, the protocol reverts back to controlled flooding.

COFSNET+ utilizes a bit field in the header to record the nodes that have rebroadcasted the packet between the source and destination. This informa-

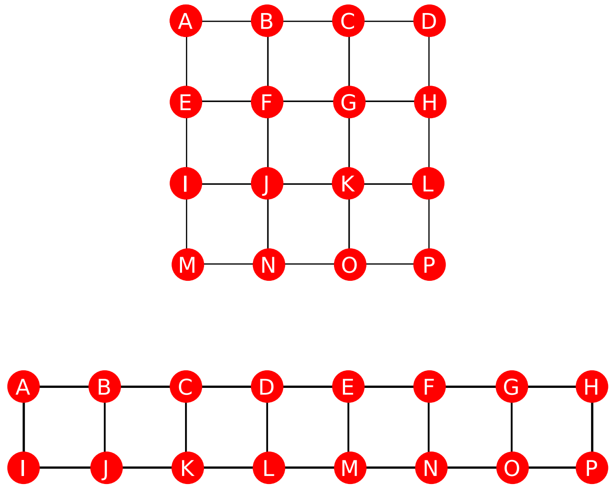

Figure 8: Grid and Ladder network topologies used in the COFSNET+ experiments.

tion can be utilized both by the intermediate nodes and by the destination node to discover the topology of the network. Furthermore, since a note may observe multiple copies of the same packet delivered over different paths, latency-based ordering among those paths can be established according to the time at which they were received.

Learned network topology information can then be used to limit which nodes rebroadcast a message, by explicitly specifying them in the outgoing packet's header. Conceptually, this represents the chain of shortest-path nodes surrounded by a band of "backup" nodes. In the most minimal approach, only the nodes on the shortest path – discovered during the previous flooding-based delivery – are included in the list. In this case, the end result would be similar to a shortest-path routing protocol such as DSR. However, the reliability of the delivery can be incrementally enhanced by including additional nodes until the desired degree of redundancy is achieved.

Because the retransmission list is only a set of nodes, it is not an exact sequential hop list. For example, consider the 4×4 Grid topology in Figure 8 and a packet sent from node A to node P. Retransmission list A,B,F,G,K,L would result in delivery along one of the shortest paths. Adding more nodes into the list (e.g., E,J,O) would increase the total number of transmissions required to deliver a packet, but it would also increase the robustness of the protocol since copies of the packet take distinct routes and some loss would be tolerated. A retransmission list containing all but the destination would effectively reduce COFSNET+ to the original COFSNET with full flooding.

Rather than outline a complete solution, COFSNET+ is meant to provide a mechanism for optimization of controlled flooding. Finding the *optimal* content of the retransmission set for a given source/destination pair – balancing a tradeoff between the reliability of COFSNET and the efficiency of shortest-path routing – is still an open question. The possible inputs include the network topology, mission information, cooperative behavior, and past reliability observations.

8 Performance Evaluation of Improved Network Protocols

8.1 Results of Packet Deprecation Simulation

Two scenarios were used to exercise the Replace and Hold method of queueing. The Cooperative Survey scenario is a simulation of the AUV Fest 2005 test run for the SAUV and the MARV where the MARV spent significant time queueing packets while disconnected from the network.

The Chain scenario was undertaken to examine the effects of using "Replace" in a network scenario that previously had been used to examine the effectiveness of AUSNET. Within this scenario, the network remains connected constantly.

Each of these scenarios were simulated multiple times, adjusting for various AUSNET settings and altering the state of the various DSR-based configuration parameters for broken route recognition.

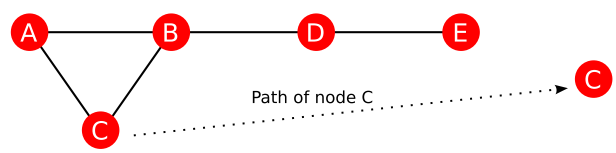

Figure 9: Chain scenario: the path travelled by node C causes changes in routes, but no breaks in connectivity.

The problem of Packet Queueing, discussed in Section 6.1, was successfully recreated within the TSI Simulator described in Section 4.1 using conditions to mimic the original run, as depicted in Figure 6. In the simulated scenario, node C moves in a large $6,000 \times 2,000$ meter simulated box while nodes A and B remain stationary inside the box. For approximately half of the time, node C is unable to communicate with nodes A or B, but is attempting to send a status packet to node A every thirty seconds. Note that the default AUSNET settings cause the protocol to retransmit a packet at most 5 times; in the presented experiments this value is varied between 2 and 9.

Use of the Replace method for packets in the queue caused a significant reduction in utilized bandwidth over the use of the Hold method. Specifically, with 64-byte status packets as shown in Figure 10, the standard 5-retry setting for AUSNET yields a 48% improvement.

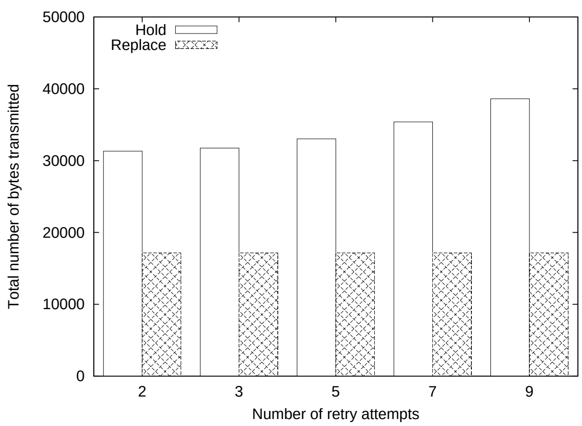

Figure 10: Simulation results for Replace-Hold scenario in Figure 6 with 64-byte status packets.

In the 5 node Chain scenario shown in Figure 9, node C is in contact with nodes A and B then travels slowly towards the endpoint of the arrow, attempting to transmit a simulated status packet every 30 seconds. Along its travel, node C will break connections, forcing packets to be enqueued before the route is rediscovered.

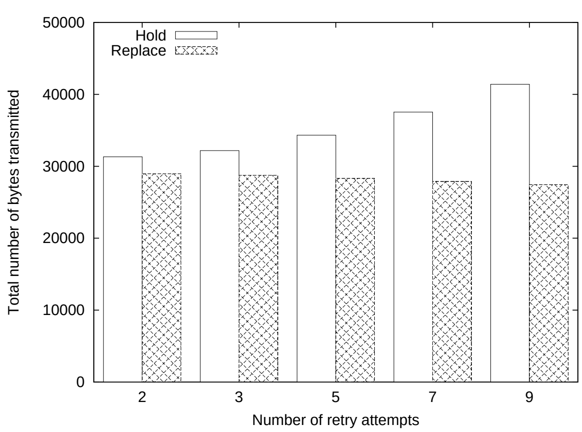

Figure 11: Results for Replace-Hold scenario in Figure 9 with 32-byte status packets.

As Figure 11 shows, the use of the Replace method yields a significant 18% improvement in network bandwidth utilization with standard 5-retry settings and 32-byte data packets.

The introduction of the Replace method into AUSNET packet queuing led to significant reduction of bandwidth usage in both scenarios. This is interesting because the algorithm was designed simply to eliminate the barrage of transmission of outdated information that would occur upon reconnection to the network. This improvement occurred in the fully-connected Chain scenario because the pace at which the replaceable data was generated exceeded the time required to determine a route to be broken. It was further shown that as the time to determine a broken route increases, the inherent value of the Replace method increases as well.

In summary, the implementation of the Replace flag for simulated status packet transmissions from AUVs showed marked improvement not only in the case of node isolation within a disjoint network, but also for normal operation in a connected network.

8.2 Results of COFSNET+ Simulation

COFSNET+ has been evaluated using the discrete event simulator described in Section 4.2, with a stream of packets sent at regular intervals. The inter-packet gap was set to a value larger than the maximum network latency to avoid packet reordering issues. An ideal, collision-free media access control (MAC) protocol was assumed.

Two regular 16-node network topologies were examined: a 4×4 *Grid* and a 2×8 *Ladder*, shown together with node labels in Figure 8. The distance between the closest nodes was set to 1 km and the transmission range was set to 1.2 km, making the diagonal nodes out of range of each other. For the Grid topology, the test traffic was sent from node A to node P, for the Ladder topology, the traffic flowed from node A to node H.

A new measure of performance was collected during these simulations that was not considered in the first development cycle: path loss. Path loss measures the probability that a packet fails to reach its intended destination. This is dependent on the individual probabilities of link loss: packet loss during transmission between a pair of nodes in the network path. In this series of experiments, the probability of link loss was varied uniformly between 0 and 1.

The performance of COFSNET+ is compared against two baseline solutions: a source-routed shortest-path protocol (e.g., DSR [5]) and controlled flooding (COFSNET). The goal is to show that COFSNET+ preserves the resiliency inherent in flooding-based protocols while reducing the overhead to a level comparable with standard ad hoc routing protocols.

A summary of the experiments for the Grid and Ladder topologies is given in Table 2 and Table 3 respectively, which assume that the network topology has been at least partially discovered. The various COFSNET+ experiments were designed to show the trade-off between the multiplicity of paths and the number of retransmissions. The tables also give the number of times a packet is retransmitted during its delivery for each of the experimental setups. For each topology and retransmission list, a single simulation was run during which 300,000 packets were generated.

Figures 12 and 13 show the packet loss and latency performance for the Grid topology. Correspondingly, Figures 14 and 15 show the equivalent results for the Ladder topology. In both cases, the presence of alternative paths between the source and destination leads to significant improvements in robustness of the packet delivery. It can be seen that the benefits of multiple paths diminish as nodes that are increasingly further away from the optimal path are added to the retransmission list. This reaffirms the design goal of allowing exploitation of the trade-off between reliability and efficiency.

The latency measurements show that at first, the average packet delivery latency increases; some of the packets are delivered over longer sub-optimal routes due to losses on the shortest path. However, for higher link loss rates, the latency – somewhat counter-intuitively – decreases. At this level of link loss, a majority of packets do not get delivered at all (see Figures 12 and 14). However, the remaining packets are likely to have been delivered over paths with fewer hops where they are targeted for loss fewer times.

9 Conclusions

Although AUSNET's proposed Replace and Hold method reduces some of the queueing problems experienced in the field, COFSNET+ appears to be more capable; we believe that it will scale from several nodes to several tens of nodes (i.e., being useful for at least the foreseeable future). As such, COFSNET+ is currently being implemented on the Solar AUV platform and readied for in-water testing.

The widely varying behaviors between AUSNET and COFSNET during in-water tests were just as much a measurement of our understanding of the physical environment as they were a measurement of

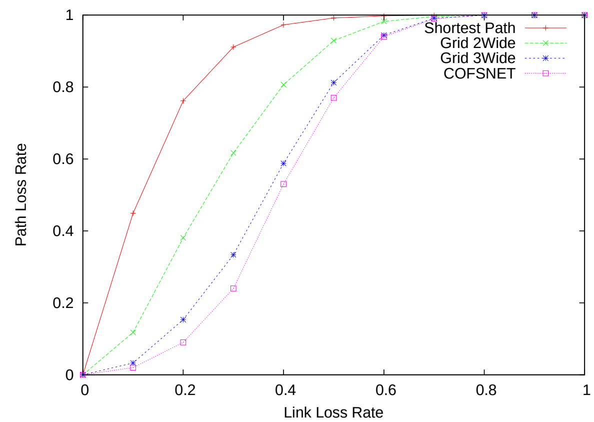

Figure 12: Packet loss probability under the Grid topology.

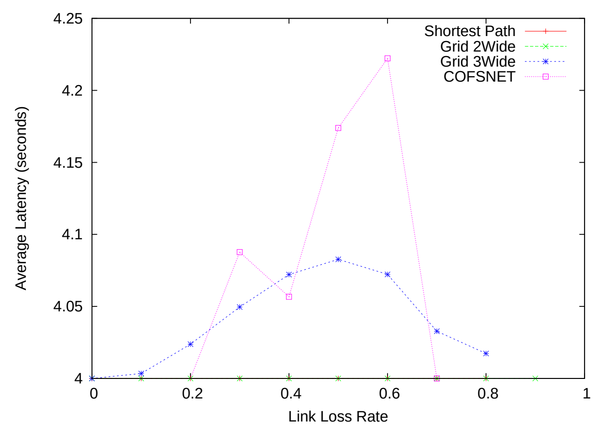

Figure 13: Packet latency under the Grid topology.

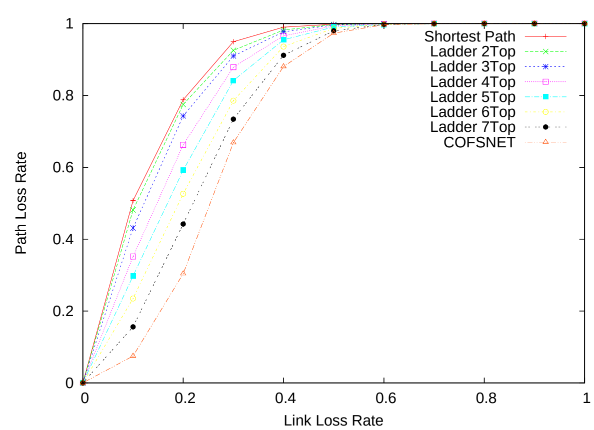

Figure 14: Packet loss probability under the Ladder topology.

Table 2: Retransmission lists for the Grid topology experiments (source A, destination P).

Protocol	Experiment	Retransmission List	No. of retransm.
Shortest Path	–	Shortest path nodes (A,B,F,G,K,L) retransmit	6
COFSNET+	2Wide	A,B,E,F,G,J,K,L,O	9
COFSNET+	3Wide	A,B,C,E,F,G,H,I,J,K,L,N,O (all but D & M)	13
Orig. COFSNET	–	All nodes except for destination retransmit	15

Table 3: Retransmission lists for the Ladder topology experiments (source A, destination H).

Protocol	Experiment	Retransmission List	No. of retransm.
Shortest Path	–	Shortest path nodes (A,B,C,D,E,F,G) retransmit	7
COFSNET+	2Top	A,B,C,D,E,F,G,I,J	9
COFSNET+	3Top	A,B,C,D,E,F,G,I,J,K	10
COFSNET+	4Top	A,B,C,D,E,F,G,I,J,K,L	11
COFSNET+	5Top	A,B,C,D,E,F,G,I,J,K,L,M	12
COFSNET+	6Top	A,B,C,D,E,F,G,I,J,K,L,M,N	13
COFSNET+	7Top	A,B,C,D,E,F,G,I,J,K,L,M,N,O	14
Orig. COFSNET	–	All nodes except for destination retransmit	15

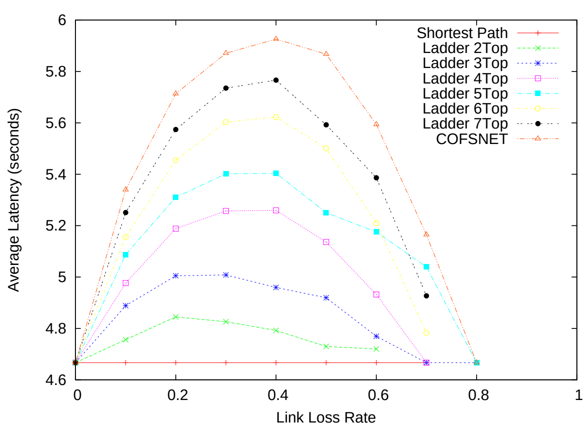

Figure 15: Packet latency under the Ladder topology.

our protocols' performance. Our experience suggests that starting with simplicity and robustness before adding efficiency may be a better path to developing viable protocols for mobile underwater acoustic networks. In a broader sense, a deeper understanding of the limitations on connectivity in a real-world environment is required before any successful protocol optimization can take place.

By failing to accurately model the random packet loss, asymmetric links, and node disconnections, our simulators did not present an accurate picture of how our protocols might behave in a harsh real-world environment. Design goals such as low overhead and scalability, present in the development AUSNET, mattered very little when a packet could not be delivered successfully in a network of only 4 nodes. The robustness of COFSNET – normally a "best effort" service – under these conditions suggests that we need to check our assumptions as to what defines a good protocol for underwater networking.

In addition to deriving and implementing more accurate models of the physical medium, simulator testing should include more measurements that reflect the practicality of a protocol. Besides the effects of individual link losses and asymmetries on path loss, overhead, and overall latency measurements, the time required for a network to recognize a route to a new node is a very important indicator of a successful protocol; nodes may enter communication range for only brief periods, requiring more immediate reactions.

This is also the point in time at which protocol designs need to be more tightly coupled with the state of the art in hardware. For example, support for one-way ranging or multiple access is not yet available in commodity acoustic modems, and there are certainly very few (if any) such devices currently deployed; protocols that make use of these features cannot undergo the scrutiny of field testing, much less be used in production code. Collaboration between the developers of these hardware and software components will enable modem designers to better meet the assumptions made by protocol designers, and protocol designers to work within the technical constraints of their access to the physical medium.

10 Future Work

In the immediate future, COFSNET+ must undergo in-water testing, both to measure the effectiveness of the protocol itself – using COFSNET as a baseline – and to verify the validity of our simulations of its performance under loss.

CADCON is the most likely candidate for simulator improvements, since it already models both vehicle movement and acoustic propagation in 3-dimensional space. Its acoustic channel model must be updated to simulate packet collisions and link asymmetry, as well as to include controls for dropping or scrambling packets. To better support post-run analysis, CADCON's logging facility must be upgraded to better record the simulated acoustic network's activity.

Due to the severely constrained aspects of the underwater communication environment, all resources available to the vehicles should be explored to determine which, if any, can contribute to a more efficient or robust design. Can the vehicle's navigation system be used to predict areas of future connectivity? Can the mission planner decide which communications are most important to the overall goal of the AUV deployment? Can protocol layers interact in a meaningful way, such that both improve their performance? Can a disruption-tolerant behavior be used alongside current "best effort" protocols? Can the vehicle itself be used to physically carry large quantities of data between points? Will more organized cooperative behaviors lead to more reliable networks? Should predictable motion patterns be used to increase opportunities for communication? The difficulty of communication demands that we use all available assistance, and the answers to these questions will help light the way to the technologies that can provide it.

Acknowledgment

This work was funded in part by the US Office of Naval Research (ONR) under grants N00014-04-1-0264, N00014-03-C-0109, N00014-02-1-008, N00014-05-1-0666 and by the National Science Foundation under grants IIS-0329837 and STTR Phase 2B program Grant DMI-032084. The authors would like to thank Aditi Goel, Aiswariya Kannan, and Jordan Thayer, graduate students of the Department of Computer Science at the University of New Hampshire, for their contributions to the development of the discrete event simulator.

References

[1] J. C. Jalbert, J. Baker, J. Duchesney, P. Pietryka, W. Dalton, D. R. Blidberg, S. G. Chappell, R. Nitzel, and K. Holappa, "Solar-powered autonomous underwater vehicle development," in *Thirteenth International Symposium on Unmanned Untethered Submersible Technology (UUST'03)*, Durham, NH, Aug. 2003.

[2] C. Benton, J. Kenney, R. Nitzel, D. R. Blidberg, S. G. Chappell, and S. Mupparapu, "Autonomous systems network (AUSNET) – protocols to support ad-hoc AUV communications," in *Proc. of IEEE/OES AUV2004: A Workshop on Multiple Autonomous Underwater Vehicle Operations*, Sebasco Estates, ME, June 2004.

[3] C. Benton, M. M. M. Haag, R. Nitzel, D. R. Blidberg, S. G. Chappell, and S. Mupparapu, "Advancements within the AUSNET protocol," in *Fourteenth International Symposium on Unmanned Untethered Submersible Technology (UUST'05)*, Durham, NH, Aug. 2005.

[4] S. S. Mupparapu, R. Bartoš, and M. M. Haag, "Performance evaluation of ad hoc protocols for underwater networks," in *Fourteenth International Symposium on Unmanned Untethered Submersible Technology (UUST'05)*, Durham, NH, Aug. 2005.

[5] D. Johnson, Y. Hu, and D. Maltz, "The dynamic source routing protocol (DSR) for mobile ad hoc networks for IPv4," RFC 4728 (Experimental), Feb. 2007. [Online]. Available: http://www.ietf.org/rfc/rfc4728.txt

[6] E. M. Sozer, M. Stojanovic, and J. G. Proakis, "Underwater acoustic networks," *IEEE Journal*

of Oceanic Engineering, vol. 25, no. 1, pp. 72–83, Jan. 2000.

[7] I. F. Akyildiz, D. Pompili, and T. Melodia, "Underwater acoustic sensor networks: research challenges," *Ad Hoc Networks*, vol. 3, no. 3, pp. 257–279, May 2005.

[8] J.-H. Cui, J. Kong, M. Gerla, and S. Zhou, "The challenges of building mobile underwater wireless networks for aquatic applications," *IEEE Network*, vol. 20, no. 3, pp. 12–18, May-June 2006.

[9] J. Partan, J. Kurose, and B. N. Levine, "A Survey of Practical Issues in Underwater Networks," in *WUWNet '06: Proceedings of the 1st ACM international workshop on Underwater networks*. Los Angeles, CA: ACM Press, 2006, pp. 17–24.

[10] P. Casari, M. Stojanovic, and M. Zorzi, "Exploiting the bandwidth-distance relationship in underwater acoustic networks," in *IEEE/MTS Oceans 2007*, Vancouver, BC, Canada, Oct. 2007, pp. 1–6.

[11] V. Rodoplu and M. K. Park, "An energy-efficient MAC protocol for underwater wireless acoustic networks," in *Proc. of IEEE/MTS Oceans 2005*, Washington, DC, Sept. 2005.

[12] X. Guo, M. R. Frater, and M. J. Ryan, "A propagation-delay-tolerant collision avoidance protocol for underwater acoustic sensor networks," in *Proc. of MTS/IEEE Oceans06 Asia-Pacific*, Singapore, May 2006.

[13] B. Peleato and M. Stojanovic, "A MAC protocol for ad-hoc underwater acoustic sensor networks," in *First ACM International Workshop on UnderWater Networks (WUWNet)*, Los Angeles, CA, Sept. 2006.

[14] M. Molins and M. Stojanovic, "Slotted FAMA: a MAC protocol for underwater acoustic networks," in *Proc. of IEEE/MTS Oceans 2006*, Singapore, May 2007.

[15] J. Yackoski and C.-C. Shen, "A frugal time-division MAC protocol for underwater acoustic networks," in *Second ACM International Workshop on UnderWater Networks (WUWNet)*, Montreal, Canada, Sept. 2007.

[16] K. B. Kredo and P. Mohapatra, "A hybrid medium access control protocol for underwater wireless networks," in *Second ACM International Workshop on UnderWater Networks (WUWNet)*, Montreal, Canada, Sept. 2007.

[17] G. G. Xie and J. H. Gibson, "A network layer protocol for UANs to address propagation delay induced performance limitations," in *Proc. of IEEE/MTS Oceans 2001*, Honolulu, HI, Nov. 2001.

[18] D. Pompili, T. Melodia, and I. F. Akyildiz, "Routing algorithms for delay-insensitive and delay-sensitive applications in underwater sensor networks," in *MobiCom '06: Proceedings of the 12th annual international conference on Mobile computing and networking*. Los Angeles, CA: ACM, 2006, pp. 298–309.

[19] I. F. Akyildiz, D. Pompili, and T. Melodia, "State of the art in protocol research for underwater acoustic sensor networks," *Mobile Computing and Communications Review*, vol. 11, no. 4, pp. 11–22, 2006.

[20] E. A. Carlson, P.-P. Beaujean, and E. An, "Location-aware routing protocol for underwater acoustic networks," in *Proc. of IEEE/MTS Oceans 2006*, Boston, MA, Sept. 2006.

[21] A. F. Harris and M. Zorzi, "On the design of energy-efficient routing protocols in underwater networks," in *Fourth Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks, 2007 (SECON '07)*, San Diego, CA, June 2007, pp. 80–90.

[22] J. A. Rice, "US Navy Seaweb development," in *WUWNet '07: Proceedings of the second workshop on Underwater networks*. Montreal, Quebec, Canada: ACM, 2007, pp. 3–4.

[23] M. Grund, L. Freitag, J. Preisig, and K. Ball, "The PLUSNet underwater communications system: Acoustic telemetry for undersea surveillance," in *IEEE/MTS Oceans 2006*, Boston, MA, Sept. 2006.

[24] T. Tether, "Statement of the director of the Defense Advanced Research Projects Agency submitted to the Subcommittee on Terrorism, Unconventional Threats and Capabilities House Armed Services Committee United States House of Representatives," Mar. 2008. [Online]. Available: http://www.darpa.mil/body/news/2008/hasc3-13-08.pdf

[25] M. M. Haag, E. Agu, R. Komerska, S. G. Chappell, and R. Bartoš, "Status packet deprecation and store-forward routing in AUSNET," in *First ACM International Workshop on UnderWater*

Networks (WUWNet), Los Angeles, CA, Sept. 2006.

[26] R. J. Komerska and S. G. Chappell, "A simulation environment for testing and evaluating multiple cooperating solar-powered AUVs," in *MTS/IEEE Oceans 2006*, Boston, MA, Sept. 2006.

[27] J. C. Jalbert, "Multiple AUV communications test report - Lake George, NY; October 17 - 22, 2004," Autonomous Undersea Systems Institute, Tech. Rep. 0411-01, Nov. 2004.

[28] ——, "SAUV II AUVFest 2005 report; june 6 – 17, 2005," Autonomous Undersea Systems Institute, Tech. Rep. 0507-02, July 2005.

[29] S. G. Chappell, R. J. Komerska, D. R. Blidberg, C. N. Duarte, G. R. Martel, D. M. Crimmins, M. A. Beliard, R. Nitzel, J. C. Jalbert, and R. Bartoš, "Recent field experience with multiple cooperating solar-powered AUVs," in *Fifteenth International Symposium on Unmanned Untethered Submersible Technology (UUST'07)*, Durham, NH, Aug. 2007.

[30] I. Katz, "A delay-tolerant networking framework for mobile underwater acoustic networks," in *Fifteenth International Symposium on Unmanned Untethered Submersible Technology (UUST'07)*, Durham, NH, Aug. 2007.

[31] R. Bartoš, R. Komerska, S. Chappell, and M. M. Haag, "COFSNET+: a routing protocol for communication among autonomous underwater vehicles," in *Fifteenth International Symposium on Unmanned Untethered Submersible Technology (UUST'07)*, Durham, NH, Aug. 2007.

Authors

Radim Bartoš received a Ph.D. degree in Mathematics and Computer Science from the University of Denver in 1997. He is currently an Associate Professor at the Department of Computer Science of the University of New Hampshire. Besides underwater networking, his research interests include access networks and theory of networked computation.

Steven G. Chappell earned a M.S. in Computer Science at the University of New Hampshire in 1987. His thesis was implementation of a prototype AUV mission planner based on a dual blackboard architecture. He is currently a Software Systems Engineer at the Autonomous Undersea Systems Institute (AUSI); formerly the Marine Systems Engineering Laboratory (MSEL). Since joining AUSI/MSEL in 1981, he has been designing, implementing, and maintaining AUSI's AUV software.

Rick J. Komerska earned his M.S. in Civil Engineering (Transportation Systems) from the University of California, College of Engineering in 1995. While at AUSI, he worked on various research and development efforts focused around multiple cooperating AUVs. Prior to AUSI, he worked as a research scientist in the Data Visualization Research Lab at the Center for Coastal and Ocean Mapping, University of New Hampshire as well as a software engineer at Cabletron Systems as a key developer in designing a performance collection architecture for a large network monitoring application.

Matthew M. Haag earned a B.S. in Computer Science at the Worcester Polytechnic Institute in 2005. His Senior Research involved communication between partitioned underwater networks using AUSNET. He was employed by Technology Systems, Inc. (TSI) from 2004 until 2007 engaging in research and development of underwater networking protocols and AUV Mission planning. He continues to support these ends as a private consultant for TSI.

Sai S. Mupparapu earned a M.S. in Electrical Engineering from the University of New Hampshire in 2003. His thesis work focused on developing a Robust and Energy Efficient Protocol for Ad Hoc Networks. Prior to joining AUSI as Research Project Engineer in March 2003, Sai worked as a Research Assistant at the University of New Hampshire InterOperability Laboratory (UNH-IOL) from 2000 – 2002. In 1999 he worked as an Assistant Systems Engineer for Tata Consulting Services in Madras, India.

Emmanuel Agu is currently an associate professor in the Computer Science Department at Worcester Polytechnic Institute. Dr. Agu received M.S. and Ph.D. degrees in Electrical and Computer Engineering from the University of Massachusetts at Amherst in 1996 and 2001 respectively. His research interests are in mobile graphics, mobile and ubiquitous computing, wireless networking and computer graphics. He has authored or co-authored over 30 publications in these areas. His research has been funded by

the NSF, the US Department of Education and AMD Inc.

Ian Katz received a B.S. degree in Computer Science from the University of New Hampshire in 2007, where he is currently completing a M.S. in Computer Science. He has designed and maintained software for testing networking technology at the University of New Hampshire InterOperability Laboratory (UNH-IOL) since 2001 and conducted network research with AUSI since 2006.