

Optimizing Network Performance through Packet Fragmentation in Multi-hop Underwater Communications

Stefano Basagni*, Chiara Petrioli†, Roberto Petroccia† and Milica Stojanovic*

* ECE Department

Northeastern University

{basagni, millitsa}@ece.neu.edu

† Dipartimento di Informatica

Università di Roma “La Sapienza”

{petrioli, petroccia}@di.uniroma1.it

Abstract—Random access protocols for multi-hop underwater wireless sensor networks that use control packets such as RTS/CTS have been shown to reduce or eliminate collisions between data packets, but they typically remain prone to collisions between control and data packets. Although control packets are usually much shorter than the data packets, when a control/data collision occurs, the entire data packet may have to be discarded. To reduce the effect of this type of collision, long data packet can be partitioned into smaller fragments confining the disruptive effect of a collision only to few fragments, so that only those need to be retransmitted. Despite the higher overhead (each fragment carries physical and the MAC headers and information on packet re-assembling), fragmentation reduces the overall traffic and the number of re-transmissions in the network. This technique is investigated in conjunction with the distance-aware collision avoidance protocol (DACAP). Simulation results show that data packet fragmentation offers benefits to throughput efficiency, end-to-end latency and energy consumption.

I. INTRODUCTION

Underwater wireless networking has been recognized as an enabling technology for a wide spectrum of applications that include ocean observation for scientific exploration or commercial exploitation, coastline protection, and prediction of underwater seismic and volcanic events [1]–[5]. The major challenges found in the design of underwater acoustic networks are the long propagation delay and low bandwidths. To address these issues research has been active on various topics in both deterministic and random access networks. The focus of our present work is on the latter type of network, where a number of nodes (users) access the channel in the same bandwidth, submitting their requests randomly as the demand dictates. This type of channel sharing is suitable for situations where each node’s traffic is bursty, consisting of packets that arrive at a sufficiently low rate that they do not require deterministic channel allocation (in fact, it would be wasteful to allocate the channel

to a node that is not using it). The nodes access the channel using a medium access control (MAC) protocol whose task is to reduce the number of packet collisions, and strike a balance between information throughput and energy consumption.

A number of MAC protocols have recently been proposed to specifically address the long-delay low-bandwidth problem of the acoustic channel [6]–[13]. Here, we focus on the Distance-Aware Collision Avoidance Protocol (DACAP) [9], which combines carrier sensing multiple access (CSMA) principles with medium access collision avoidance (MACA). This protocol uses dedicated control packets (request-to-send/clear-to-send (RTS/CTS) and short warning packets) to prevent collisions between data packets. However, while the absence of data/data collisions is guaranteed, control/data collisions are still possible. The effect of these collisions is often neglected, but, as we shall see, they can significantly harm the system performance. In fact, it was shown recently [14]–[16], that if the environment is characterized by frequent channels failures, in particular for Ad Hoc and resource limited sensor networks, retransmission-based protocols can lead to really long delays and possibly zero throughput channel. In a wireless setting, it is possible (and common) for two nodes that cannot directly hear each other to impair each other’s reception. Figure 1 illustrates such a situation.

This situation is exacerbated in acoustic channels, where the spreading factor (path loss exponent) is low. Through extensive simulations on most of the underwater MAC protocols proposed so far, we have observed that the vast majority of packet losses are due to this type of interference [13]. Specifically, in the case of RTS/CTS-based access à la IEEE 802.11 with the distributed coordination function, we observed that 90% of packet losses are due to interference coming from nodes that are outside of the receiver’s transmission range. This




Figure 1. Although A and B are outside of each other's nominal range, B's transmission to D will reach A, interfering with reception from C. The interference radius is thus greater than the nominal transmission radius R .

occurs even in networks where the traffic is not particularly high. Moreover, many of these collisions happen between control and data packets (65% are control/data, 10% are data/data and 15% are among control packets). Although totally eliminating control/data collisions, the use of two different channels for control and data packets [17] increases the source-to-sink packet latency and it is also ineffective in case of high channel BER.

Here we investigate the use of packet fragmentation and selective repeat ARQ at the MAC level for reducing the effect of control/data collisions. By partitioning a long packet into smaller fragments we aim at confining the disruptive effect of a collision only to few fragments, so that only those need to be retransmitted, thus reducing the overall traffic and improving the throughput efficiency, end-to-end latency and energy consumption performances.

Earlier work on this topic includes [18]–[20]. In [18] the authors enhance system goodput in CSMA/CA MAC protocol using static fragmentation. In [19], [20], two adaptive/dynamic fragmentation algorithms are proposed, both of which utilize channel measurements to dynamically select the optimal number of fragments with minimal network overhead. All these solutions are designed for RF networks, where the propagation delay is close to zero and the bit rates are on the order of Mbps. These solutions assume that after each fragment transmission the destination replies with an ACK to give an immediate feedback to the transmitter. In underwater acoustic channels, however, this is not possible due to the low speed of sound propagation. To address this problem, we investigate the performance of DACAP with and without data packet fragmentation. Our results show that despite the higher overhead (each fragment carries physical and MAC headers and information packet re-assembling), our solution yields performance improvement in terms of throughput efficiency, end-to-end latency and energy consumption.

The remainder of the paper is organized as follows. Section II describes data packet fragmentation. Perfor-

mance of DACAP, using a different number of fragments for each data packet, is discussed in Section III. Concluding remarks are given in Section IV.

II. DATA PACKET FRAGMENTATION

In [21] it was shown that considering an underwater acoustic channel almost error-free ($BER = 10^{-6}$) longer packets offer better channel utilization. In contrast to this situation, at $BER = 10^{-4}$, there clearly exists an optimal packet size for which the throughput efficiency is maximized. Assuming that each node has long data packets to send, in this section we investigate the use of packet fragmentation at the MAC level for reducing the effect of control/data collisions.

We assume that each node fragments a data packet into k fragments. Each fragment is associated with an $ID \in [1 \dots k]$. The k fragments are sent back to back in a group and a cumulative ACK is awaited in return. The receiver checks each fragment individually and for each fragment it sets the i -th bit of the ACK to 1 if the i -th fragment has been received correctly, and to 0 otherwise. Each receiver has to store the fragments until they can be rearranged before forwarding the data (Figures 2).


We consider two different ways of fragmenting packets.

- 1) Each transmitter sends out a group of k fragments and awaits the ACK, but only those fragments that are negatively acknowledged are re-transmitted in the next group, i.e., no new packet fragments are sent along with those re-transmitted.
- 2) Each transmitter sends out a group of k fragments and awaits the ACK. The fragments that are negatively acknowledged are placed in a new group of k fragments, i.e., the re-transmitted fragments are sent together with the new packet's, if any. We add two features to this scheme: a) If the number of old fragments waiting to be sent is h , with $h \leq k/5$, then a group of $k + h$ fragments is sent. This speeds-up the transmission of the entire data packet. Without this feature, the next group will miss h new fragments, thus requiring two transmissions instead of one. b) Old fragments are sent at the end of the group. The reason for this is that fragments at the beginning of the group are more likely to be affected by control/data collision and interference, due to the underwater propagation delay.

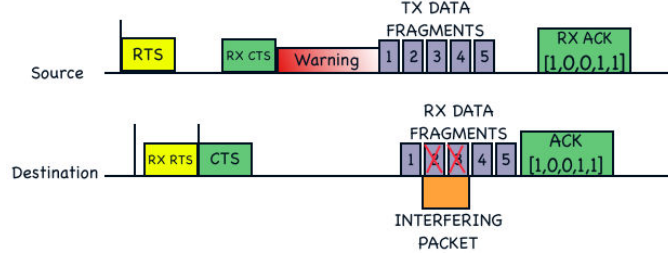
Using these schemes, we investigate the performance of DACAP. We refer to the protocol simply as DACAP if the first method of fragmentation is used, and DACAP-A if the second method is used.

III. PERFORMANCE EVALUATION

DACAP was implemented in the VINT project ns-2 simulator [22] extended to include key characteristics



(a) DACAP: All fragments are correctly delivered



(b) DACAP: Fragments 2 and 3 have to be re-transmitted

Figure 2. Fragmenting data packets in group of 5 fragments.

of the underwater environment such as 3D deployment, propagation at the speed of sound, and acoustic path loss that depends on the distance and frequency.

A. Simulation scenario

We consider a scenario with 100 static nodes placed on the seafloor at a depth of 200m. The nodes are randomly and uniformly scattered within a $4\text{km} \times 4\text{km}$ footprint. Packets are transmitted from the nodes to a common sink (data collection point) located centrally on the surface. All nodes have the same transmission radius $R = 1000\text{m}$. Communication occurs in a multi-hop fashion, over shortest path routes, which are pre-determined. The receiving power and the idle power are set to 100mW. The transmission power, as well as the carrier frequency, are computed according to [23], so as to achieve the SNR of 20dB at a distance of R . The transmission power is set to 4930mW. The nodes are equipped with half-duplex acoustic modems transmitting at a carrier frequency of 24KHz with a bandwidth efficiency of 1bps/Hz. We consider a basic (raw) transmission rate of 2000bps. The data payload size is set to 3000B, so as to optimize the throughput efficiency [21]. Two different values of BER are considered, 10^{-6} and 10^{-4} . Traffic is generated according to a Poisson process with aggregate (network-wide) rate λ packets per seconds. We also define the normalized packet rate as $\bar{\lambda} = \lambda T_{\text{pack}}$, whose values are considered in the range 0 to 1 packets per packet time. The packet time used here is the one corresponding to transmission at full rate, i.e., $T_{\text{pack}} = N_b/R_b$, where N_b is the packet size in bits and $R_b = 2000\text{bps}$. Simulations

were conducted ranging from very low traffic ($\bar{\lambda} = 0.01$), to medium-high traffic ($\bar{\lambda} = 0.3$).

Once a packet is generated it is associated with a source selected randomly among all the nodes. The destination of all the packets is the sink. The total size of the data packet is given by the payload plus the headers added by different layers (physical through network). The physical layer header contains all the information needed by the modem to start receiving a packet (synchronization preamble, delimiters, etc.). A synchronization peering time is taken to be 10ms. We consider the cases with no fragmentation, and different numbers of fragments per packet. The number of fragments ranges from 5 to 100, in increments of 5. Each fragment carries physical and MAC headers and information on packet re-assembling. The MAC header contains the sender's ID, the destination's ID and the packet type if needed. The re-assembling information added to each fragment is set to 3B and it includes the remaining number of fragments in the group that the transmitter is sending, the ID of the fragment within the data packet, and the total number of fragments for the data packet (in this way it is possible to vary the number of fragments for each data packet). The MAC header length is set to 3B, while the ACK length is set to 3B plus the number of bytes containing the information to ack or nack the fragments that have been sent. The size of the RTS and CTS packets is set to 6B each. The size of each fragment is set to the number of bytes containing N_b/k bits, where k is the number of fragments into which each

data packet is partitioned.

To correctly receive each packet (control or data) the signal to interference ratio at the receiver is required to be $SIR \geq 15\text{dB}$. Each node has a buffer of 30KB where data coming from the upper layers are stored before transmission. Whenever the buffer is full and a new packet arrives, the oldest packet is discarded. We also limit the number of packets that can be stored to 50, so that the nodes are not filling their buffers with old information.

Our implementation of DACAP mandates to abandon RTS transmission after 7 failed attempts to access the channel, and to discard a data packet after 7 failed re-transmissions. Every point depicted in our figures was obtained by averaging over the number of simulations needed to achieve a statistical confidence of 95% with a 5% precision.

B. Performance metrics

The following metrics have been used to assess the system performance.

1) *Throughput efficiency*, defined as the ratio between the bit rate delivered to the sink (correct bits) and the bit rate offered to the network, $N_b\lambda$.

2) *Energy per bit*, defined as the energy consumed by the network to correctly deliver a bit of data to the sink.

3) *End-to-end latency per meter*, defined as the time between packet generation and the time when it is correctly received by the sink, divided by the distance between the source and destination. Normalization by the distance is used to unify the performance over a varying coverage area (a large area will entail proportionately large propagation delay). This metric is computed only for the packets correctly delivered, and averaged over all the packets.

4) *Route length*, defined as the average number of hops traversed by packets correctly delivered to the sink.

In what follows we discuss these performance measures as functions of the offered load and different BERs.

C. Simulation results

1) *Throughput efficiency* is shown in Figure 3. This figure indicates the advantage of fragmenting long packets. We can see how using at least two fragments increases the number of correctly delivered packets. Increasing the number of fragments increases the overhead, and for this reason, there exists an optimal number of fragments from the viewpoint of throughput efficiency. This optimal number is listed in Table I for two values of the BER.

Figure 3(a) shows that when the traffic load increases and the number of control/data collisions becomes higher, DACAP without fragmentation (first point in the curves) exhibits a throughput efficiency of 0.4, while

Table I
OPTIMAL NUMBER OF FRAGMENT FOR 3000B LONG DATA PACKETS.

	BER= 10^{-6}	BER= 10^{-4}
DACAP	15	30

DACAP with fragmentation and DACAP-A achieve an efficiency of 0.65 and 0.7, respectively. DACAP-A performs slightly better than DACAP (with fragmentation) because it reduces the total number of channel accesses.

Figure 3(b) shows remarkable improvement in data delivery at all traffic loads. There exists a smallest number of fragments for which the performance can be brought to its best. More importantly, we note that this number, which is on the order of 30 for the example considered, is practically manageable.

2) *Energy consumption per bit* is shown in Figure 4. These results reveal another advantage of fragmentation: Reducing the effect of control/data collisions reduces the overall energy consumption, despite the fact that a higher overhead is introduced for each data transmission. The greatest advantage comes at high traffic and higher BER. In these cases, DACAP saturates and the throughput efficiency decreases, but the energy wasted is considerably less than in the case of no fragmentation. At the lower BER, DACAP saves about 25% of energy using 15 fragments instead of sending the entire packet. At the higher BER, the use of more than one fragment results in much higher energy savings. The best results are achieved, as said before, using about 30 fragments.

3) *Packet latency per meter* is shown in Figure 5. Differently from multiplexing control and data channels [17], Figure 5(a) shows how the use of packet fragmentation can reduce the packet latency. Packet latency per meter, especially at higher traffic load, is much lower. Hence, choosing the number of fragments close to the optimum, higher throughput efficiency, lower energy consumption, and shorter packet latency are obtained simultaneously (Table I). Note that for long-range transmissions, a small reduction in latency for each meter can reduce the overall end-to-end packet latency by minutes.

Considering a number of fragments ranging from 10 to 15 at medium-high traffic load (Figure 5(a)), DACAP-A shows a packet latency per meter of about 0.5s while that of DACAP without fragmentation is 1s, i.e., without fragmentation the end-to-end latency is twice longer. At high BER, without fragmentation, many packets are lost due to errors. As we can see from Figure 6(b), with no fragmentation only nodes close to the sink are able to correctly deliver data. This means that communications coming from nodes far away in the network are lost and never make it to the sink, reducing the traffic load in the

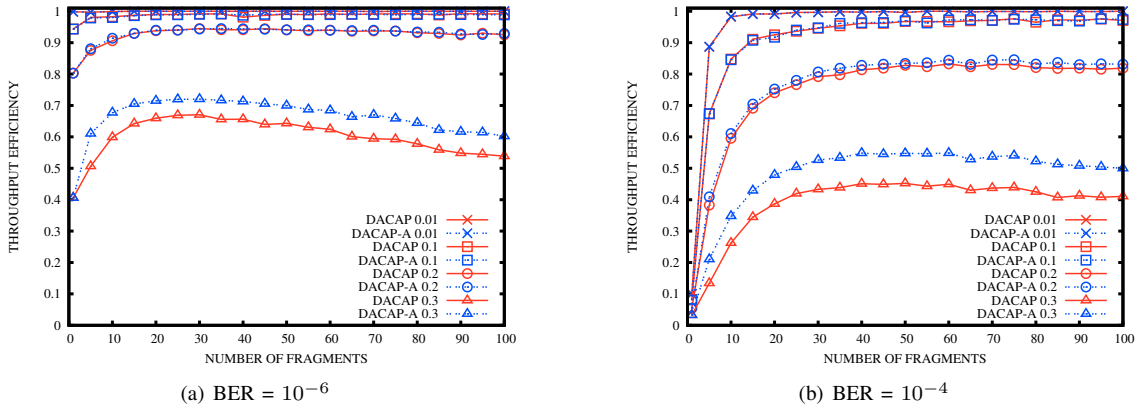


Figure 3. Throughput efficiency.

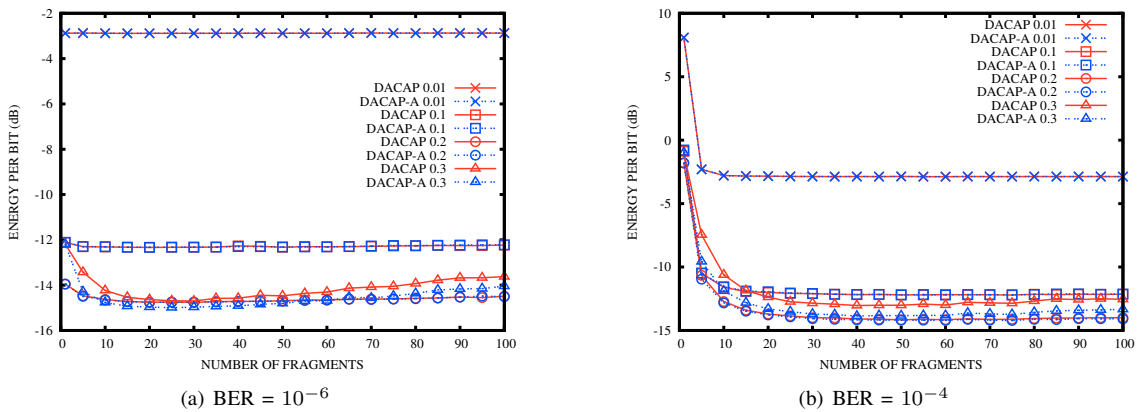


Figure 4. Energy consumption per bit.




Figure 5. Packet latency per meter.

sink area and reducing the packet latency. In comparison, fragmentation considerably reduces the BER effect and increases the “fairness.”

4) *Route length* is shown in Figure 6. The average number of hops is seen to decrease with traffic. This means that as the number of re-transmissions increases and the network becomes congested, nodes closer to the

sink are those that deliver more packets successfully. This is confirmed by the fact that a node is on the average 2.35 hops away from the sink, but that the average number of hops traveled by a successful packet, when the traffic is high, is less than this quantity. Packet fragmentation significantly reduces this effect allowing nodes far away in the network to correctly deliver data

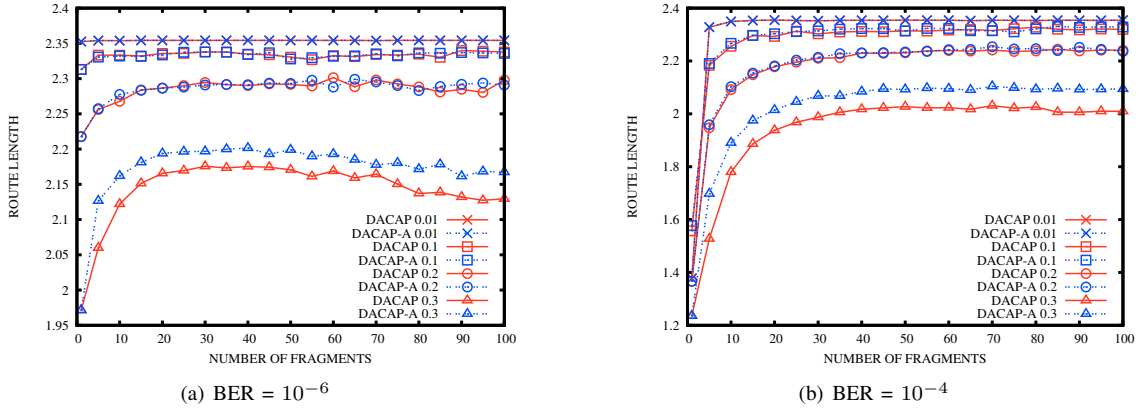


Figure 6. Length of the routes traveled by packets correctly delivered to the sink.

to the sink avoiding unfairness on packet delivery. This effect is particularly pronounced in case of high BER (Figure 6(b)).

IV. CONCLUSIONS

The use of packet fragmentation and selective repeat ARQ was considered for random access underwater networks that use MAC protocols based on RTS/CTS exchange. We focused on a particular MAC protocol, the delay-aware collision avoidance protocol (DACAP). The goal in doing so was to confine the disruptive effect of a collision to only few fragments, so that only those need to be re-transmitted, thus reducing the overall traffic. Fragmentation was shown to increase throughput efficiency, while simultaneously reducing end-to-end latency and energy per bit consumption. The benefits of fragmentation are particularly pronounced at low BER values, which are typical of highly distorted underwater acoustic channels.

ACKNOWLEDGMENTS

This work was supported in part by the EU FP 7 STREP project CLAM “CoLIAborative EMbedded Networks for Submarine Surveillance” and by the NSF grants 0946610 and 0831728.

REFERENCES

- [1] I. F. Akyildiz, D. Pompili, and T. Melodia, “Underwater acoustic sensor networks: Research challenges,” *Elsevier’s Journal of Ad Hoc Networks*, vol. 3, no. 3, pp. 257–279, May 2005.
- [2] J.-H. Cui, J. Kong, G. M., and Z. S., “Challenges: Building scalable mobile underwater wireless sensor networks for aquatic applications,” *IEEE Network, Special Issue on Wireless Sensor Networking*, vol. 20, no. 3, pp. 12–18, May 2006.
- [3] J. Heidemann, W. Ye, J. Willis, A. A. Syed, and Y. Li, “Research challenges and applications for underwater sensor networking,” in *Proceedings of the IEEE Wireless Communications and Networking Conference, WCNC 2006*, Las Vegas, NV, April 3–6 2006, pp. 229–235.
- [4] 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*. New York, NY, USA: ACM, September 2006, pp. 17–24.
- [5] L. Lanbo, Z. Shengli, and C. Jun-Hong, “Prospects and problems of wireless communication for underwater sensor networks,” *Wireless Communications and Mobile Computing, Special Issue on Underwater Sensor Networks*, vol. 8, no. 8, pp. 977–994, August 2008.
- [6] J. Yackoski and C. Shen, “UW-FLASHR: Achieving high channel utilization in a time-based acoustic MAC protocol,” in *WuWNeT ’08: Proceedings of the third ACM international workshop on Wireless network testbeds, experimental evaluation and characterization*, San Francisco, California, USA, September 15 2008, pp. 59–66.
- [7] N. Chirdchoo, W.-S. Soh, and K. C. Chua, “ALOHA-based MAC protocols with collision avoidance for underwater acoustic networks,” in *Proceedings of IEEE INFOCOM 2007, the 26th IEEE International Conference on Computer Communications*, Anchorage, Alaska, USA, May 6–12 2007.
- [8] M. Molins and M. Stojanovic, “Slotted FAMA: A MAC protocol for underwater acoustic networks,” in *Proceedings of MTS/IEEE OCEANS 2006*, 2006.
- [9] B. Peleato and M. Stojanovic, “Distance aware collision avoidance protocol for ad-hoc underwater acoustic sensor networks,” *IEEE Communications Letters*, vol. 11, no. 12, pp. 1025–1027, December 2007.
- [10] V. Rodoplu and M. Park, “An energy-efficient MAC protocol for underwater wireless acoustic networks,” in *Proceedings of MTS/IEEE OCEANS 2005*, 2005.
- [11] X. Guo, M. Frater, and M. Ryan, “An adaptive propagation-delay-tolerant MAC protocol for underwater acoustic sensor networks,” *IEEE Journal of Oceanic Engineering*, vol. 34, no. 2, pp. 170–180, April 2009.
- [12] A. Syed, W. Ye, and J. Heidemann, “Comparison and evaluation of the T-Lohi MAC for underwater acoustic sensor networks,” *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 9, pp. 1731–1743, December 2008.
- [13] C. Petrioli, R. Petroccia, and M. Stojanovic, “A comparative performance evaluation of MAC protocols for underwater sensor networks,” in *Proceedings of MTS/IEEE OCEANS 2008*, Quebec City, Quebec, Canada, September 15–18 2008.
- [14] P. Jelenkovic and J. Tan, “Can retransmissions of superexponential documents cause subexponential delays?” in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, Anchorage, Alaska, USA, May, 6–12 2007.
- [15] —, “Characterizing heavy-tailed distributions induced by retransmissions,” in *Technical report, Department of Electrical Engineering, Columbia University, EE2007-09-07*, September 2007.

- [16] —, “Is ALOHA causing power law delays?” in *ITC 2007. 20th International Teletraffic Congress*, Ottawa, Canada, June, 17–21 2007.
- [17] S. Basagni, C. Petrioli, R. Petroccia, and M. Stojanovic, “Assessing the advantages of channel multiplexing for underwater networking,” in *Proceedings of MTS/IEEE OCEANS 2009*, Biloxi, Mississippi, USA, October, October, 26–29 2009.
- [18] Y. Chang, C. P. Lee, and J. A. Copeland, “Goodput optimization in CSMA/CA wireless networks,” in *Broadband Communications, Networks and Systems, 2007. BROADNETS 2007. Fourth International Conference on*, Raleigh, North Carolina, USA, September, 10–14 2007.
- [19] P. R. Jelenković and J. Tan, “Dynamic packet fragmentation for wireless channels with failures,” in *MobiHoc '08: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking and computing*. Hong Kong, Hong Kong, China: ACM, May, 27–30 2008, pp. 73–82.
- [20] Y. Chang, C. Lee, B. Kwon, and J. Copeland, “Dynamic optimal fragmentation for goodput enhancement in WLANs,” in *Testbeds and Research Infrastructure for the Development of Networks and Communities, 2007. TridentCom 2007. 3rd International Conference on*, Orlando, Florida, USA, May, 21–23 2007.
- [21] S. Basagni, C. Petrioli, R. Petroccia, and M. Stojanovic, “Choosing the packet size in multi-hop underwater networks,” in *Proceedings of IEEE OCEANS 2010*, Sydney, Australia, May, 24–27 2010.
- [22] The VINT Project, *The ns Manual*. <http://www.isi.edu/nsnam/ns/>, 2002.
- [23] M. Stojanovic, “On the relationship between capacity and distance in an underwater acoustic communication channel,” *SIG-MOBILE Mob. Comput. Commun. Rev.*, vol. 11, no. 4, pp. 34–43, 2007.