A Hybrid Medium Access Control Protocol for Underwater Wireless Networks

Kurtis B. Kredo II
Electrical and Computer Engineering Dept.
University of California, Davis
kbkredo@ucdavis.edu

Prasant Mohapatra
Computer Science Department
University of California, Davis
prasant@cs.ucdavis.edu

ABSTRACT

Underwater networks allow investigation of many areas of the world not easily accessed by humans, but offer interesting challenges to device and protocol designers due to the unique channel conditions present when using acoustic communications. The high transmit power of acoustic transceivers makes the medium access protocol a primary focus point for reducing energy consumption in energy limited underwater devices. Scheduled protocols use very little power by eliminating collisions, but cannot adapt to changing traffic conditions in the same way as random protocols. We attempt to bridge these two ideas by dividing time into scheduled and unscheduled access periods in order to yield the benefits of each protocol. We show that this technique increases the bits delivered per energy unit in many cases of interest. Additionally, the hybrid technique provides low latency for a wider range of traffic rates than either of the two protocols when considered individually. We also investigate some of the design tradeoffs to consider when using a hybrid protocol.

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocols

General Terms: Design, Performance

1. INTRODUCTION

Underwater wireless networks enable researchers to monitor many areas of interest that are difficult for humans to directly access. However, the use of acoustic communication in underwater environments presents many interesting challenges to protocol and device designers [12, 1]. Low data rate and large propagation delay are common traits of underwater acoustic communication. Typical acoustic channel data rates are on the order of several kilobits per second, as compared to the megabits per second common in radio frequency (RF) wireless communication. Acoustic signal propagation occurs at approximately 1500 meters per second, several orders of magnitude slower than the 3×10^8 meters per second experienced in RF propagation.

WUWNet'07, September 14, 2007, Montréal, Québec, Canada.

Underwater transceivers also have larger power ratios than in other networks. Table 1 lists the approximate power consumption for devices typical in underwater (WHOI modem), RF sensor (Mica 2), and RF wireless computer networks (Cisco Aironet) [5, 3]. Underwater transceivers have transmit powers orders of magnitude higher than RF devices and a higher ratio of transmit to receive power, so protocols that utilize the acoustic radio effectively become much more important in underwater networks. These constraints are further complicated by the limited energy resources available to underwater devices when powered by batteries.

Table 1: Node Power Consumption

State	Underwater	RF Sensor	RF Computer
Tx	50 W	80 mW	2.24 W
Rx	3 W	30 mW	1.35 W
Idle	80 mW	30 mW	1.35 W

Medium access control (MAC) protocols directly impact how and when nodes utilize their wireless transceiver. MAC protocols must balance the need to conserve energy, and thus extend the network lifetime, with application requirements, which may change over the network's lifetime. One method to decrease energy consumption, by avoiding collisions, is to use a time division multiple access (TDMA) protocol. However, TDMA cannot easily adapt to changing traffic conditions. Alternatively, random access MAC protocols can quickly adapt to traffic conditions, but suffer from energy waste through packet collisions.

One approach to reduce energy and maintain adaptability involves combining these techniques. This work attempts to utilize TDMA and an unscheduled channel access method in a single, hybrid protocol to provide low energy consumption through reduced collisions while still adapting to changing traffic conditions. We do this by dividing time into a TDMA portion, which has no collisions and low energy consumption, and an unscheduled access portion, which adapts to changing traffic conditions. Additionally, nodes utilize state information, such as queue depth or number of contenders, shared during the TDMA portion to better divide channel access during the unscheduled portion. We show that combining TDMA, state distribution, and an unscheduled protocol reduces the energy required to deliver packets in many cases. Additionally, we find that a hybrid protocol can deliver packets at a low, constant latency for a wider range of data rates than either protocol individually. However, our results indicate that MAC protocols must use the state information effectively to yield improved performance.

We present related work proposed in the literature in the next section and then discuss our proposed protocol in Section 3. Section 4 presents analysis of the proposed protocol for various parameter choices and presents an initial investigation into the technique. Further depth is provided through simulation, as presented in Section 5. We conclude in Section 6.

2. RELATED WORK

Several MAC protocols have been proposed recently that attempt to provide sufficient operation despite the unique conditions present in underwater acoustic networks. Most proposals have focused on random access techniques similar to MACA [6] to reduce the number of collisions, but some have used a fully synchronized approach.

The Slotted FAMA [7] protocol applies a slotted structure to the Floor Acquisition Multiple Access (FAMA) [4] protocol proposed for RF networks. Similar to MACA, nodes coordinate communication through the use of request to send (RTS) and clear to send (CTS) packets. However, in Slotted FAMA, nodes can only send these control packets at the beginning of a slot. With the proper selection of slot size, the protocol prevents data packet collisions. Similarly, Tone Lohi [15], uses tones, instead of control packets, to contend for the wireless channel and estimate channel contention. Our protocol does not use control packets to coordinate channel access, but instead relies upon state information included in earlier packets.

Peleato and Stojanovic [8] adapt the MACA protocol to underwater networks by adjusting the time required for control packet exchange based on the distance between the sender and receiver, which decreases the overhead associated with sending data packets. After receiving a CTS, the source node waits for a certain time, calculated based on the round trip time measured in the RTS/CTS exchange, before sending the data packet. The source and destination use the delay period and short warning messages to prevent data packet collisions that may arise from neighboring transmissions. This protocol reduces the overhead associated with Slotted FAMA by reducing control packet exchanges, but may still yield a large number of control packet collisions and large backoff latencies in dense networks.

Chirdchoo et al [2] proposed several adaptations to the Aloha protocol that leverage the large delay present in underwater acoustic communications. Aloha with advanced notification (Aloha-AN) nodes transmit a short notification packet before transmitting the much longer data packet after some delay. The notification packets allow nodes to collect more information about an intended destination and limit data packet collisions. Our work takes a more progressive approach to reduce collisions by distributing state reliably during TDMA time periods. Additionally, as the traffic rate increases, a purely random approach will increase the number of collisions and decrease performance.

Rodoplu and Park [10] present a protocol that coordinates communication without using a slotted structure. In their protocol, each node randomly selects a time to send synchronization messages, which contain the time period between transmissions. The synchronization packet sizes are kept small in comparison to the transmission period (less than 1%) to limit collisions. After developing a schedule that separates node transmission times, nodes sleep between their neighbors' wakeup time and their selected wakeup time. The protocol relies upon very low duty cycles to reduce the number of collisions, so it may not be applicable to applications that require higher data rates. Our work attempts to reduce collisions, even during periods of high activity, while maintaining low energy operation when the network has little traffic.

Another protocol [11] uses a combined TDMA and CDMA approach based on clustering. Within each cluster, nodes communicate using a TDMA schedule setup by the cluster head and a CDMA code assigned to that cluster. Intercluster communication occurs through nodes within range of multiple clusters (the authors assume nodes have multiple packet reception). Our technique also uses TDMA, but introduces adaptability to changes in traffic conditions. Our current work only considers a single, fully connected network, so the the proposed ideas could be applied to our work to enable inter-cluster packet forwarding.

TRAMA [9], designed for terrestrial sensor networks, operates similarly to our proposed protocol. Sensor nodes using TRAMA limit data packet collisions by organizing channel access based on shared topology and traffic information. A TRAMA frame consists of several random access signaling slots, which are used to share topology and traffic information, followed by organized access transmission slots used for data transmissions. After determining a schedule based on the traffic information, sensor nodes either receive packets destined for them or sleep during each transmission slot. Our protocol shares the state information during scheduled slots, thus reducing the number of collisions associated with state distribution. We also propose resource division algorithms which are much less computationally intensive for the nodes.

All of the previous protocols attempt to balance the energy spent through protocol overhead and the energy saved by reducing data packet collisions. Random access contention schemes, such as tones or RTS/CTS exchanges, invest energy and channel utilization in control packets in an effort to decrease data packet losses. These schemes are more useful when data packets are much larger than control packets. Scheduled protocols, whether TDMA or a slotted structure, likewise invest energy and channel utilization through coordination and synchronization packets to prevent or reduce collisions. Which scheme performs best is heavily dependent on the network topology, data patterns, and other aspects. TDMA based MAC protocols offer reduced energy consumption, especially in heavily congested networks, since collisions are eliminated, but lose their advantage in networks with low traffic volume, when the traffic is unequally distributed among the nodes, or when traffic conditions change frequently. Random MAC protocols have an advantage in low congestion, but waste energy through collisions as the traffic intensity increases.

3. HYBRID MAC PROTOCOL

In an effort to combine the best of both channel access schemes, we propose a hybrid protocol that includes scheduled and unscheduled periods in a slotted frame. The scheduled portion of the frame allows nodes to communicate without collision and guarantees a certain data rate available to each node in the network, while the unscheduled portion allows nodes to adapt to changing traffic conditions. Additionally, the scheduled portion allows nodes to distribute state information quickly and reliably.

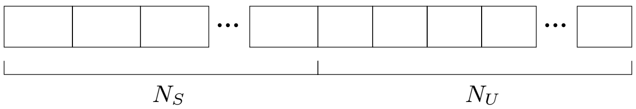

Figure 1: Hybrid Protocol Frame Structure

We assume that all nodes can overhear each other and listen to the channel at all times. For example, the nodes may be in a cluster within a larger network using CDMA to limit interference between clusters or the nodes may be part of a small network centered around a buoy gateway. Future work may include expanding the protocol to handle multihop networks and allowing nodes to sleep to conserve energy if appropriate [5]. The focus of this work is to examine hybrid MAC protocol schemes that operate between the two extremes of scheduled and random access, and to highlight some of the design issues associated with a hybrid protocol.

Figure 1 shows the frame structure for the protocol. The first portion of the frame is divided into N_S scheduled slots using TDMA, where each node is assigned one slot for transmission. The number of scheduled slots can be defined by the network operator or a clustering protocol, but is assumed to be constant once network operation begins. After the scheduled slots, are N_U unscheduled slots. We use the term unscheduled slots because the access mechanism for these slots may vary. For example, nodes may compete for the slots through a random access protocol or the slots may be temporarily assigned to nodes based on the distributed state. The main distinction is that each scheduled slot is assigned to one node for a very long time, while the unscheduled slots will be used by various nodes in different frames. We evaluate a few example unscheduled slot access mechanisms in the next sections, but many others are possible. One way to assign the unscheduled slots would involve the use of a centralized controller, but in this work we assume no such controller is present and base our designs on a distributed approach.

Each time slot is long enough to transmit a maximum length packet plus the longest expected propagation delay. Adding the propagation delay to the slot time ensures that nodes will completely receive a packet before another node begins transmitting. The long slot times create the disadvantage that nodes remain idle for long periods of time, but this work does not focus on reducing time slot length for several reasons. First, a shorter time slot would cause collisions between nodes, even when they transmit in separate time slots, and would cause collisions between nodes transmitting in scheduled slots and unscheduled slots. Additionally, shorter slots may still allow the central destination node used in this work (see later sections) to receive packets without collisions, but it would decrease the reliability of state information distribution as collisions may still occur at other nodes. This would be particularly bad as nodes may then divide the unscheduled slots inappropriately, causing many further collisions. A balance exists between reducing the idle energy wasted in long time slots and the additional collisions experienced when using short time slots. Exploring this balance has been left for future work.

Using a slotted structure requires nodes to maintain synchronization with their neighbors. This is a particularly difficult problem in underwater acoustic networks due to the time varying channel conditions and is currently being investigated by others [14]. We assume that nodes, through the use of higher quality clocks and synchronization protocols, can maintain slot synchronization with their neighbors. The long slot times (on the order of several seconds) minimize the effect of clock drift and synchronization inaccuracy, but buffer time may be added to prevent collisions.

The operation of a hybrid MAC protocol will heavily depend upon the protocols chosen for operation during each segment and the amount of each frame assigned to the protocols. Additionally, any information shared between the segments, such as the state information shared during the scheduled access portion, can have a large impact on protocol operation.

4. ANALYSIS

We now examine a hybrid MAC protocol at a high level, as much of the specifics depend on the access mechanism chosen for the unscheduled slots. In the following analysis we assume each node with data to send is backlogged (i.e., each node either has no packets to send or infinitely many).

First, consider the energy spent by the network during operation. Define E_{RX} as the energy expended by a node during a slot in which it receives a packet, E_I as the energy expended during a slot while idle, and E_{TX} as the energy expended during a slot in which a node transmits a packet. The transmission and reception energy may not remain constant across slots, as packet sizes vary, so we define E_{TX} and E_{RX} as the averages taken over all slots where a transmission or reception occurs. Furthermore, define T as the number of users with packets to send and N as the total number of users in the network. The average energy consumption for a TDMA protocol is then

$$E_T = \frac{TE_{TX} + (N-1)TE_{RX} + N(N_S - T)E_I}{N_S} \quad (1)$$

The energy expended by an unscheduled access protocol will depend on the actual method used, but will be at least

$$E_U = \frac{P_R N_U E_{TX} + (N-1)P_R N_U E_{RX} + N(1 - P_R N_U)E_I}{N_U} \quad (2)$$

where P_R is the probability of successful reception in a slot. Collisions and transmitting any control packets will cause the unscheduled protocol's average energy to increase. For the hybrid protocol we get a combination of the two, for an average energy of

$$E_H = \frac{E_T N_S + E_U N_U}{N_S + N_U} \quad (3)$$

Thus, the new protocol will have an average energy between a TDMA protocol and the access mechanism chosen for the unscheduled portion of the frame. Figure 2 shows the energy consumption of the protocols with various parameter values as P_R varies. Two points are worth noting. First, changing the value of T only changes the operational range of the hybrid protocol, the slope does not change. This can be attributed to the fact that increasing T increases the energy consumed in the TDMA portion of the frame, which does not depend on P_R. (For now we ignore that P_R will vary with T in the unscheduled access protocol.) Second, adjusting the value of N_U changes the slope of the energy curve (where N_S is held constant at 10 in the graph). As N_U increases, the hybrid protocol operates more like the unscheduled access

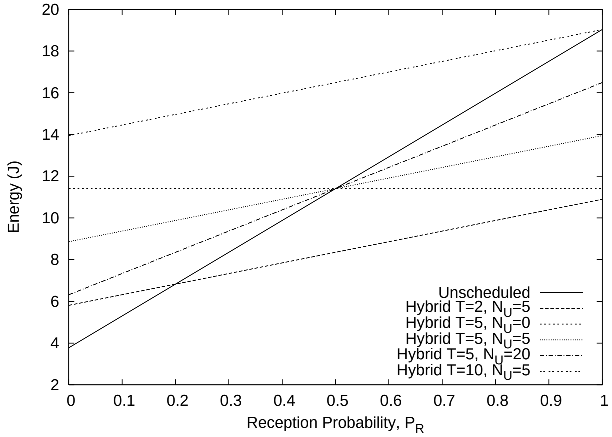

Figure 2: Protocol Energy

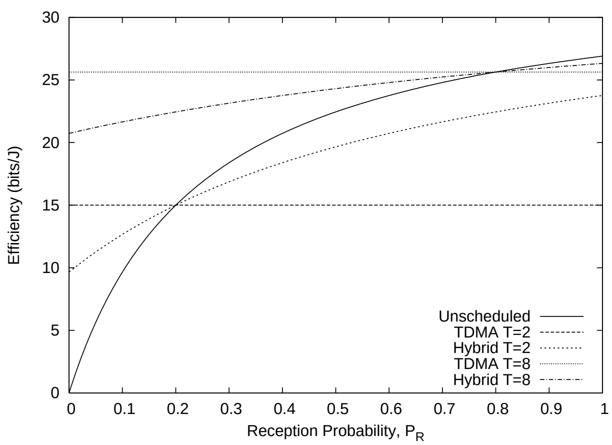

Figure 3: Abstract Protocol Efficiency

protocol, so the energy curve slope increases. The reverse is true as N_U decreases. This shows how a hybrid protocol can be used to balance node energy consumption with traffic conditions by varying the ratio of N_S to N_U.

Energy, by itself, is not a good evaluation metric since the minimal energy would have no nodes transmit any packets. A better metric is to consider the amount of benefit provided by the network for the energy expended. We define efficiency as the number of upper layer bits delivered divided by the total energy expended in the network. Figure 3 shows the efficiency of the protocols (using 512 bit packets) for several operational points as P_R varies. Notice that the unscheduled protocol curve intersects the TDMA curve when $\frac{T}{N_S} = P_R$ and that the hybrid protocol is not the optimal solution for any particular value of P_R, but it does better than either individual protocol over some span of P_R.

For a more concrete evaluation, we compare four simple MAC protocols: TDMA, a simple random protocol, a mixture of the previous two, and a protocol, called divider, that assigns the unscheduled slots based on state information shared during the scheduled slots. Nodes using the random protocol transmit in each slot with probability P_T when they have data to send, and was chosen for a simple analysis. When the power consumed by a node is constant and

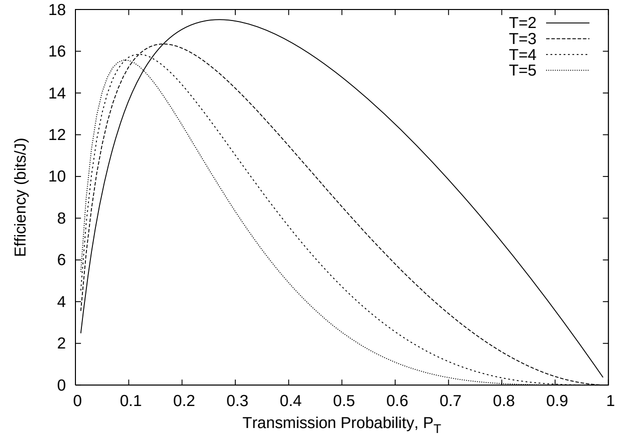

Figure 4: Random Protocol Efficiency

independent of its operating state, the optimal P_T equals $\frac{1}{T}$. However, this is not the case in underwater networks. The efficiency of the simple random protocol is easily found as

$$\mathcal{E} = \frac{B(1, T, P_T)}{B(0, T, P_T) N E_I + \overline{E_1}} \tag{4}$$

where

$$\overline{E_1} = \sum_{i=1}^{T} B(i, T, P_T)\left(i E_{TX} + (N - i) E_{RX}\right)$$

and $B(x; t, p)$ is the binomial probability distribution function on variable x over t experiments with success probability p. $\overline{E_1}$ equals the average energy consumption in each slot when at least one node decides to transmit. Figure 4 plots the efficiency of the random protocol for several values of T. Notice that, since $E_{TX} > E_{RX} > E_I$ (see Table 1), the optimal probabilities have decreased from $\frac{1}{T}$ to avoid the energy losses caused by collisions. Also note that as T increases, the added channel congestion causes the achievable efficiency to decreases.

Figure 5 shows the efficiency of all the protocols as P_T varies for T equal to two and eight. $N = N_S = N_U = 10$ in this example. When T equals two, we see that the random protocol can achieve a higher efficiency than TDMA for the proper choice of P_T. However, when T is greater than 2, no value of P_T allows the random protocol to achieve a higher efficiency. For all operating points the hybrid protocol operates between the TDMA and random protocols, as expected. The divider protocol achieves a very high efficiency since it can fully utilize the unscheduled slots without collisions. Figure 5 further illustrates that the hybrid MAC protocol performance will depend on the unscheduled access mechanism.

The analysis thus far has shown the benefits and weaknesses of using scheduled and unscheduled access protocols individually, and the capabilities of using a hybrid protocol to combine the individual benefits. However, it has made several unrealistic assumptions to ease in analysis. First, channel errors, ignored in the analysis, will reduce the number of packets delivered. This will lower the utilization of each protocol, but may be particularly damaging for protocols, such as the divider protocol, that distribute state. Any packet lost during a scheduled slot will result in nodes

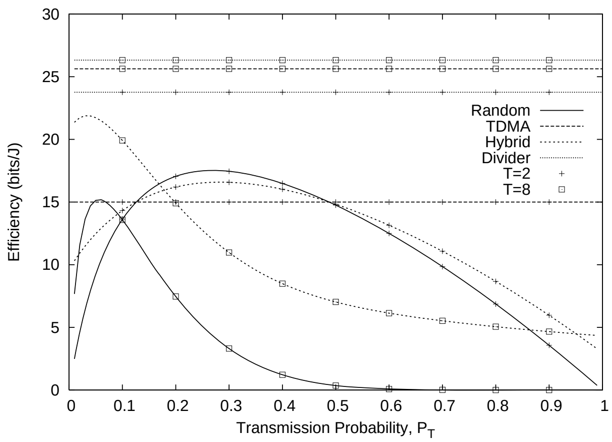

Figure 5: Concrete Protocol Efficiency

having incorrect state and may increase collisions. Second, the amount of traffic generated at each node will vary over time. Nodes may not always have traffic to send, as the above analysis assumed. This would impact the simple random protocol described earlier as T will vary across frames, so a single, optimal value for P_T will not exist. The traffic dynamics also affect the values chosen for N_U. A very large N_U may seem like a good way to increase efficiency when using the divider protocol, but this will delay traffic generated at nodes until the next scheduled portion of the frame. For these reasons, and to gain a better understanding of the protocol dynamics, we explore the ideas further in simulation.

5. SIMULATION RESULTS

In the following simulations we model each node as a stationary device that communicates acoustically. Each node transmits at a fixed power of 120 dB re μPa using binary phase shift keying (BPSK) at a rate of 2500 bps. The power consumed in the transmit, receive, and idle state are 40 W, 3 W, and 80 mW, respectively. The network topology consists of a ring of $N = 10$ nodes, which generate traffic according to a Poisson process, evenly distributed around a central receiver, which is the destination of all packets. The topology might be used when several nodes communicate with a surface buoy used to relay information to a remote collection point. The acoustic channel is modeled using practical spreading (a path loss exponent of 1.5) and common modeling equations [13]. Each node communicates at a frequency of 5 kHz using a bandwidth of 5 kHz. The channel noise is set to 40 dB re μPa. Under the given channel and node parameters, nodes can communicate at a distance of 6735 m with a signal to noise ratio of 20 dB. Channel losses are modeled based on BPSK error equations scaled so an SNR value of 20 dB corresponds to a bit error rate of 10^{-3}. Each simulation was performed 20 times with different random inputs and the average results are reported with 95% confidence intervals on the initial results, but are left off of later results for clarity. The confidence intervals are similar on all results. All packets contain 512 upper layer bits and 64 header bits (including MAC layer header fields and physical layer preamble). The Divider protocol, detailed below, adds an 8-bit field for state information. No packet retransmissions occur in any of the protocols, so any packets lost due to channel errors or collisions are silenty dropped. Each time slot is set to the sum of the packet tranmission length, the maximum propagation delay for the communication range, and a 1 millisecond guard period. For all simulations, time slots are 4.7214 seconds long. Unless stated otherwise, $N_S = N_U = 10$.

We simulated several simple protocols:

- Random: a simple random scheme where nodes transmit in a slot with a constant probability when they have packets to send; $N_S = 0$
- TDMA: a TDMA scheme where each node has a unique slot assigned during simulation initialization; $N_U = 0$
- Hybrid: a combination of the Random and TDMA schemes with $N_S = N_U = 10$

Additionally, we evaluate two protocols that distribute state information during the scheduled slots for use in organizing access to the unscheduled slots. The stateful protocols are:

- Aware Hybrid: the above Hybrid protocol where nodes adjust their transmit probability based on binary state information (whether the node needs access to the unscheduled slots); nodes use the optimal probabilities as found through simulation; nodes use a transmit probability of 0.5 when they do not detect other transmitters
- Divider: nodes include the number of packets they have queued in scheduled access packets and proportionally divide the unscheduled slots among nodes with packets to send

We first need to determine the optimal transmission probabilities for the Random and Hybrid protocols. Figure 6 shows the efficiency (bits per Joule) for the Random protocol as the transmit probability, P_T, varies. The network generates packets at a rate of 0.4 packets per second evenly divided among the transmitting nodes. Only T nodes generate traffic in these simulations since we are interested in finding the optimal transmit probability given a certain number of contenders. Note that the graph shape closely matches that suggested by the analysis, but the new channel model has resulted in slightly lower optimal transmit probabilities. Further simulations, not shown for clarity, revealed the optimal transmit probability for all values of T between two and ten, which are listed in Table 2. We use the same procedure to find the optimal transmit probabilities for the Hybrid protocol and show the results in Figure 7 and Table 2. Overall, the optimal transmit probabilities decrease from the optimal found for the Random protocol. Notice, however, that as T increases the Hybrid protocol efficiency increases due to the increase in efficiency of the TDMA portion of the frame. This is in contrast to the Random protocol, which decreases in maximum efficiency as T increases.

Using these values we evaluate the protocols by varying the traffic rate (packets per second) generated by the network, as shown in Figure 8 and 9 where $T = 2$ and $T = 8$, respectively. Note that the Hybrid protocol operates between the two individual protocols, as predicted by the analysis. When $T = 2$, the TDMA protocol operates poorly due to low channel utilization and the Random protocol operates well due to low channel contention. When $T = 8$, the

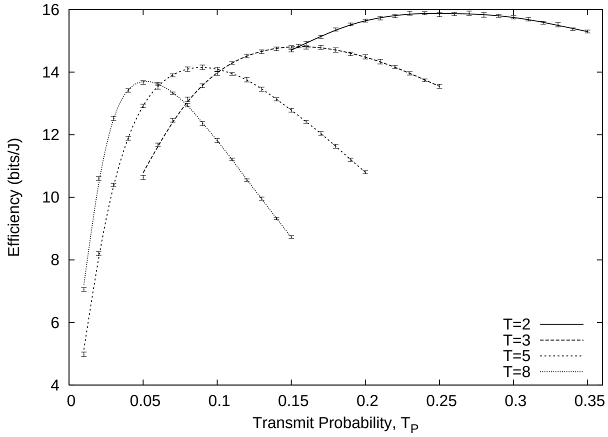

Figure 6: Random Protocol Efficiency

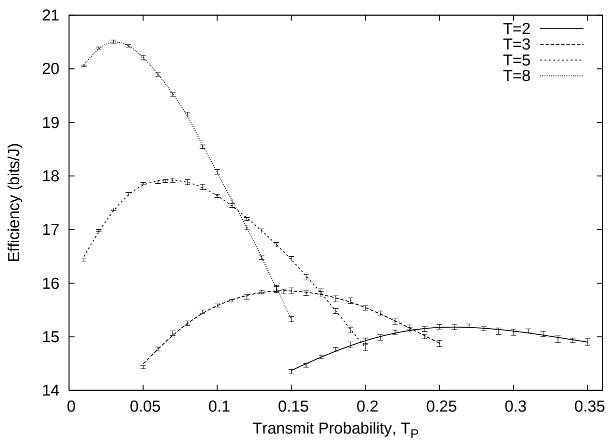

Figure 7: Hybrid Protocol Efficiency

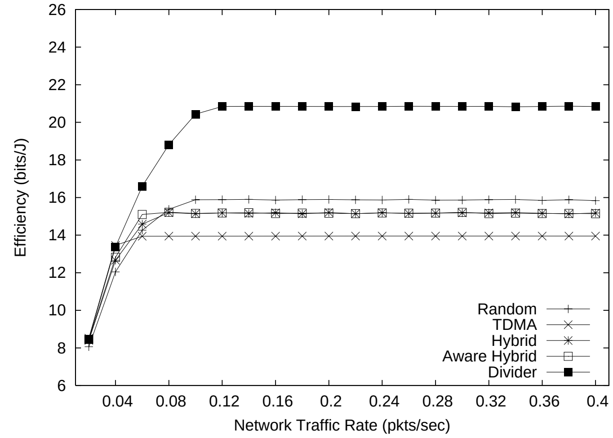

Figure 8: Protocol Efficiency for T=2

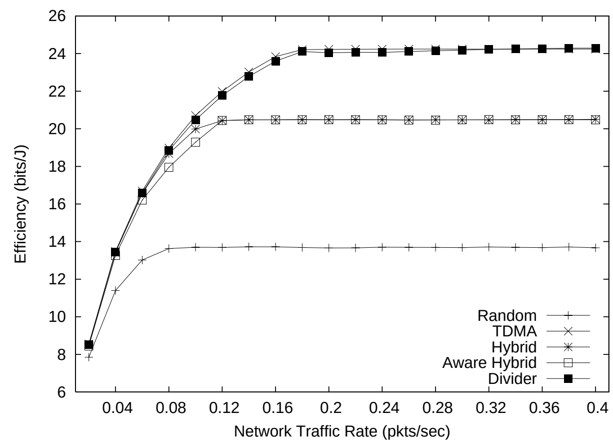

Figure 9: Protocol Efficiency for T=8

TDMA and Random protocols switch and TDMA becomes more efficient, as suggested by analysis. The Divider protocol achieves a very high efficiency under both conditions, easily beating all other protocols when $T = 2$ and closely matching TDMA when $T = 8$. The Aware Hybrid protocol closely matches the Hybrid protocol when the network is saturated (above 0.08 for $T = 2$ and 0.12 fro $T = 8$), but differs below that. When the channel contention is low, $T = 2$, the Aware Hybrid protocol gains some advantage, but when the contention increases, $T = 8$, incorrect state due to packet losses results in lower efficiency.

Figure 10 shows the efficiency of the protocols as the number of transmitters varies when the traffic created by the network is large (0.40 packets per second). As mentioned earlier, the Random protocol performs poorly and only achieves a slightly larger efficiency than the TDMA and Hybrid protocols when $T = 2$. Note also that the Random protocol uses the optimal transmit probability for each value of T, so the efficiency would be greatly decreased if a constant transmit probability were used for all cases. The TDMA protocol achieves an efficiency directly proportional to T as more time slots get used and fewer slots are wasted when all nodes remain idle. The Hybrid and Aware Hybrid protocols operate identically as the Aware Hybrid protocol cannot take advantage of variations in the number of transmitters. Over most of the values the Divider protocol operates the best. When T is greater than eight, the TDMA protocol maintains a higher efficiency than the Divider protocol by 0.9%–1.4%, but for T less than eight, the Divider protocol acheives a higher efficiency by at least 2% over the next best protocol. The Divider protocol cannot match the efficiency of TDMA in all cases due to collisions in unscheduled access slots caused by incorrect state information.

Figure 11 and Figure 12 show the message latency of the protocols as the traffic created by the network varies with $T = 2$ and $T = 8$. Message latency was measured as the time between the creation of the packet in the source node to the time the central destination receives it. Since the network does not involve routing, the latency is only impacted by the

Table 2: Optimal Transmit Probabilities

T	Random	Hybrid
2	0.25	0.26
3	0.155	0.145
4	0.11	0.095
5	0.09	0.065
6	0.07	0.05

T	Random	Hybrid
7	0.06	0.04
8	0.05	0.03
9	0.045	0.025
10	0.04	0.02

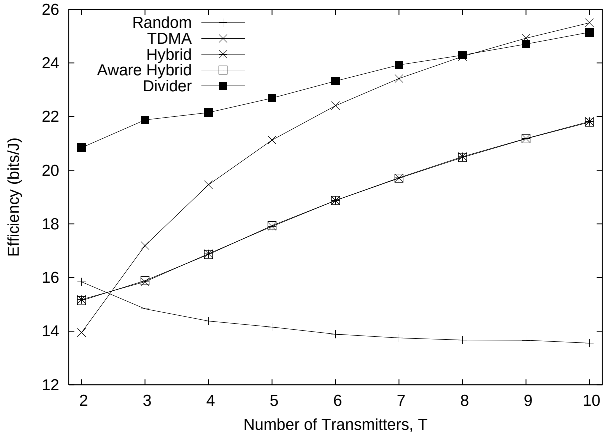

Figure 10: Protocol Efficiency as T Varies

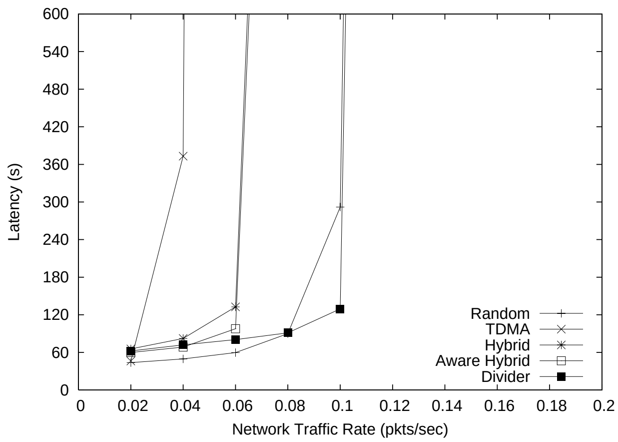

Figure 11: Message Latency for T=2

MAC protocol. With low channel contention ($T = 2$), the Random protocol can achieve a low message latency and the TDMA protocol has the highest latency. Notice also that the Divider protocol maintains a low relative latency until the network becomes congested and the latency increases to very high levels. The other protocols reach congestion at a lower rate than the Divider protocol. The same holds true for the Divider protocol when then the channel contention increases ($T = 8$). Notice also that the TDMA and Random protocols have changed positions as the Random protocol now has a large latency due to the low efficiency-optimal transmission probability. Figure 12 also shows that the Hybrid Aware protocol achieves a lower latency than the Hybrid protocol at the expense of decreased efficiency, as discussed previously. Thus, the protocols that share state information can support a higher throughput, for a similar energy consumption, than protocols that do not share state. In particular, the Divider protocol has a relatively constant delay for data rates the protocol can support and requires only slightly more energy per bit when the channel has many contenders.

One disadvantage for protocols that share state is that they might have to transmit at a higher power so nodes besides the current destination can overhear the state information (an alternative approach would have nodes distribute

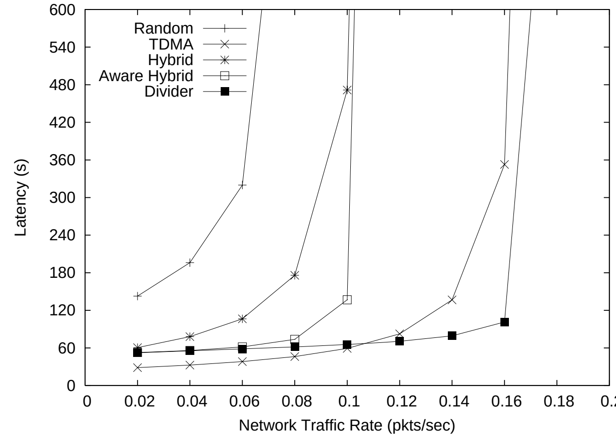

Figure 12: Message Latency for T=8

information about their neighbors, but this also involves extra energy expenditure through packet overhead). In the topology considered so far, the TDMA protocol could transmit at a lower power and still allow the central destination to receive the packets. To study the impact of lower power operation, we modified the protocols slightly by providing a low power mode where the nodes could transmit at 115 dB re μPa and consume 15 W. The TDMA protocol only uses the low power mode for transmitting. The Hybrid, Aware Hybrid, and Divider protocols transmit at high power during scheduled slots and transmit at low power during unscheduled slots. Thus, packets that have state information are transmitted at a higher power so more nodes can overhear them, but packets that don't contain relevant state information are transmitted at a lower power to save energy. Figure 13 shows the protocol efficiency as the number of transmitters vary when the traffic rate is high (0.40 packets per second). Notice the low power efficiency results are very similar to Figure 10, but the efficiency values have increased. The Divider protocol decreases in relative performance to TDMA as it now only achieves a higher efficiency when T is less than five and has a much lower efficiency than the TDMA protocol for T greater than five. The Divider protocol still has several advantages over TDMA (for example, the nearly constant latency and higher achievable throughput for a given delay), but the range where it attains a higher efficiency has decreased. One possible improvement for the Divider protocol would be to decrease the size of packets sent at a high power (in these results all packets have the same size), but this is left to future exploration.

Protocols that share state information also have the ability to adapt to changing traffic conditions. We investigate this by having nodes generate groups of packets in a burst with a certain probability. The network has all nodes ($T = N = 10$) generate traffic for the central receiver at a low rate of 0.02 messages per second. When a node decides to generate a message, it creates either 10 packets in a burst or a single packet and sends them to the MAC layer. Adaptive protocols can utilize the state information they collect to better handle the changes in traffic conditions. Figure 14 shows the efficiency of the protocols as the probability of generating a burst increases. The TDMA, Hybrid, and Divider protocols achieve very similar efficiencies for the var-

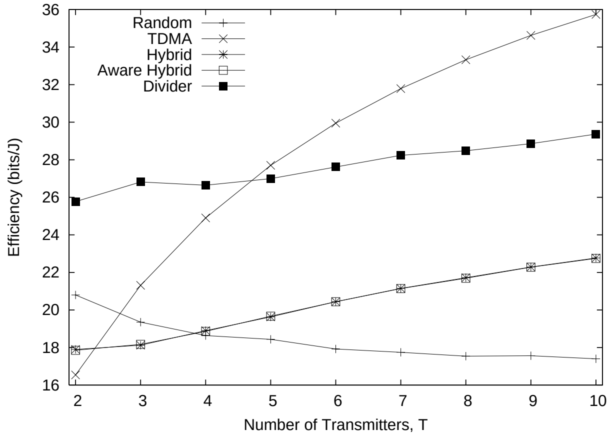

Figure 13: Low Power Protocol Efficiency

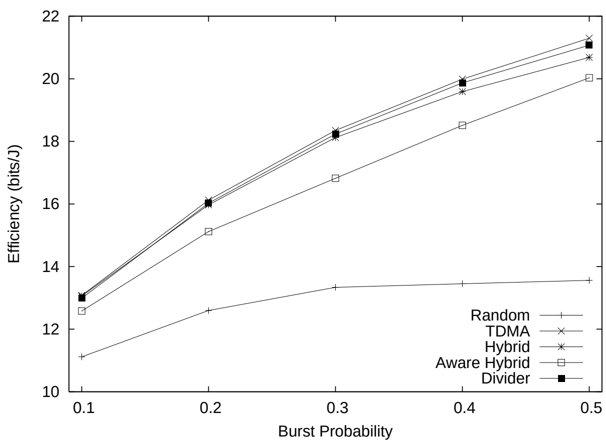

Figure 14: Protocol Efficiency for Bursty Traffic

ious burst probabilities. However, the Aware Hybrid and Random protocols achieve low efficiency due to many collisions (the Aware Hybrid protocol frequently underestimates channel contention and uses a higher transmit probability). Figure 15 shows the packet latency for the protocols and resembles previous results. Sharing state information can decrease the packet delay at the cost of slightly lower efficiency. Also, the Divider protocol can support higher burst probabilities while maintaining a lower delay as the nodes may utilize slots as needed.

6. CONCLUSION

We have shown that a hybrid protocol created by mixing a scheduled access and unscheduled access protocols can perform better than either protocol individually in certain situations. The protocol efficiency is higher in many cases, but may not be optimal for dense and heavily loaded networks as collisions cause it to waste energy when compared to TDMA. A hybrid protocol may also provide a lower, more constant latency for a wide range of traffic rates. Conversely, for a given latency, a hybrid protocol can handle a higher traffic rate. Similarly, for bursty traffic conditions, a hybrid protocol can achieve a lower latency for a very similar efficiency.

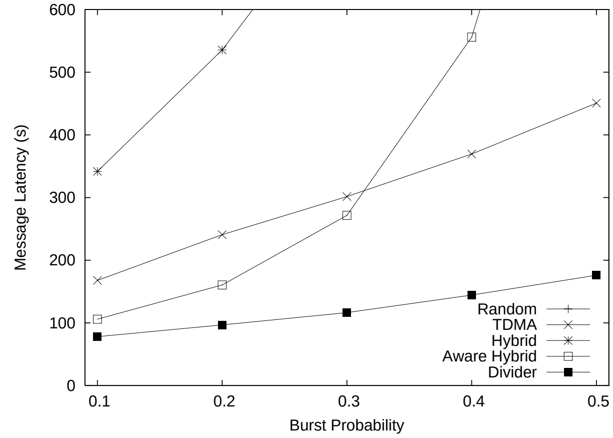

Figure 15: Protocol Latency for Bursty Traffic

This work has shown the benefits of a hybrid protocol, but also illustrated that care must be taken when choosing the protocols and the parameters of operation. For example, the Aware Hybrid protocol performed poorly when compared to the Divider protocol, illustrating that protocols must not only share state information, but use it effectively. Further work will refine the optimal settings for hybrid protocols and the conditions conductive to their efficient operation.

7. REFERENCES

[1] I. F. Akyildiz, D. Pompili, and T. Melodia. Underwater acoustic sensor networks: Research challenges. *Ad Hoc Networks*, 3(3):257–279, May 2005.

[2] N. Chirdchoo, W.-S. Soh, and K. C. Chua. Aloha-based MAC protocols with collision avoidance for underwater acoustic networks. In *Proc. IEEE InfoCom*, May 2007.

[3] Crossbow Mica2 Datasheet.

[4] C. L. Fullmer and J. J. Garcia-Luna-Aceves. Floor acquisition multiple access (FAMA) for packet-radio networks. In *Proc. ACM SIGCOMM Conf.*, pages 262–273, Aug. 1995.

[5] A. F. Harris III, M. Stojanovic, and M. Zorzi. When underwater acoustic nodes should sleep with one eye open: Idle-time power management in underwater sensor networks. In *Proc. ACM WUWNet*, pages 105–108, Sept. 2006.

[6] P. Karn. Maca - a new channel access method for packet radio. In *ARRL Computer Networking Conf.*, Sept. 1990.

[7] M. Molins and M. Stojanovic. Slotted FAMA: A MAC protocol for underwater acoustic networks. In *IEEE Oceans*, May 2006.

[8] B. Peleato and M. Stojanovic. A MAC protocol for ad-hoc underwater acoustic sensor networks. In *Proc. ACM WUWNet*, pages 113–115, Sept. 2006.

[9] V. Rajendran, K. Obraczka, and J. Garcia-Luna-Aceves. Energy-efficient, collision-free medium access control for wireless sensor networks. In *Proc. ACM SenSys Conf.*, pages 181–192, Nov. 2003.

[10] V. Rodoplu and M. K. Park. An energy-efficient MAC protocol for underwater wireless acoustic networks. In *Proc. IEEE Oceans Conf.*, volume 2, pages 1198–1203, Sept. 2005.

[11] F. Salva-Garau and M. Stojanovic. Multi-cluster protocol for ad hoc mobile underwater acoustic networks. In *Proc. IEEE Oceans Conf.*, volume 1, pages 91–98, Sept. 2003.

[12] E. M. Sozer, M. Stojanovic, and J. G. Proakis. Underwater acoustic networks. *IEEE J. Oceanic Engineering*, 25(1):72–83, Jan. 2000.

[13] M. Stojanovic. On the relationship between capacity and distance in an underwater acoustic communication channel. In *Proc. ACM WUWNet*, pages 41–47, Sept. 2006.

[14] A. A. Syed and J. Heidemann. Time synchronization for high latency acoustic networks. In *Proc. IEEE InfoCom*, Apr. 2006.

[15] A. A. Syed, W. Ye, and J. Heidemann. Medium access for underwater acoustic sensor networks. Technical Report ISI-TR-624, University of Southern California, Oct. 2006.