

Routing in Multi-Modal Underwater Networks: a Throughput-optimal Approach

Roe Diamant[‡], Paolo Casari[‡], Filippo Campagnaro[§], Michele Zorzi[§]

[‡]Department of Marine Technology, University of Haifa, Israel

[‡]IMDEA Networks Institute, Madrid, Spain

[§]Department of Information Engineering, University of Padova, Italy

Abstract—While acoustic signals are still the main communication means under water, other technologies are being developed, e.g., based on optical and radio-frequency electro-magnetic waves. Each technology has its own advantages and drawbacks to trade off, e.g., communication range vs. bit rate. Recently, new approaches are emerging to leverage the advantages of several underwater technologies by incorporating them in a multi-modal communication system. In this paper, we address a fundamental part of these multi-modal systems by proposing a novel routing protocol for networks of multi-modal nodes. Our protocol makes distributed optimal and fair decisions about the per-link flow, prevents bottlenecks, and allows simultaneous transmissions using multiple technologies in order to advance a packet towards its destination. We analyze the performance of our protocol via model-based simulations and compare it to benchmark results. Our results show that our protocol successfully leverages all technologies to deliver data, even in the presence of imperfect topology information.

Index Terms—Underwater networks; underwater acoustic communications; optical communications; multi-modal systems; optimal routing; simulations;

I. INTRODUCTION

Most physical layer (PHY) technologies developed so far to communicate under water rely on acoustic communications [1]. However, optical communications [2], [3] and radio-frequency (RF) electro-magnetic communications [4], [5] are also gaining momentum. Each technology offers a different balance of advantages and disadvantages. For instance, acoustics typically provides low (order-of-kbit/s) transmission rates, but can cover ranges up to several kilometers. However, the performance of underwater acoustics is highly influenced by the environment, especially by strong and time-varying multipath. Underwater optical communications, on the other hand, provide a very high bandwidth, up to hundreds of Mbit/s, but require the transmitter and the receiver to be aligned and only a few meters apart. In addition, optical communications are sensitive to turbidity and ambient light. RF communications do not need any alignment and can be developed based on standard terrestrial radio systems; however, the conductivity of ocean waters attenuates RF waves already after a few tens of meters [4], and limits the achievable bit rates to typically less than 100 kbit/s.

The above analysis suggests that there is a lot of potential in the integration of different PHYs into a multi-modal communication system. Such a system may be able to exploit the advantages of different technologies by transmitting through the best available one at any given moment. This approach was proposed, e.g., in [6], where the authors concluded that a system encompassing optical and acoustic communications

would be a good candidate for the wireless control of remotely operated vehicles. Recent work [7] supports the vision of multi-modal systems by showing that embedded processing platforms have become sufficiently evolved to host signal processing algorithms on general-purpose computing platforms.

The logic that decides how to switch between available PHYs plays a crucial role in multi-modal communication systems. Generally, the nodes may connect to different neighbors using (possibly partially overlapping) subsets of their PHYs. These subsets may change over time according to, e.g., environmental conditions, mobility, and traffic requirements of the nodes. In this context, the network may wish to divide the data load among different PHYs instead of always choosing links of higher capacities, which may induce bottlenecks in an upstream neighbor. In this paper, we aim to provide a multihop routing scheme that fully utilizes the available PHY technologies in an optimized fashion. To do so, we formalize the routing problem as an optimization problem where each node tries to maximize its channel goodput using all its available PHYs. Our contribution is twofold: *i*) we propose a distributed routing algorithm for multi-modal underwater networks, which maximizes the channel goodput using all technologies available to each node, while at the same time balancing the traffic flow through the network and pursuing a fair network utilization for all nodes; *ii*) we consider both the cases where incomplete and complete topology information is available to each node. We evaluate the performance of our routing algorithm by means of numerical simulations.

In the following Section II we discuss the related work; we then present and formalize our routing scheme (Section III) and describe our simulation setup and results (Section IV); Section V finally draws concluding remarks.

II. RELATED WORK

Acoustic and optical communications are typical technologies employed in multi-modal systems. A first implementation of an acoustic/optical multi-modal underwater transceiver was presented in [8], and employed for a data mule autonomous underwater vehicle (AUV). Radio-frequency technologies for underwater communications are also under development, although optical and acoustic systems still achieve higher bit rates at all distances of interest [6]. A notable feature of multi-modal systems is that the composition of multiple powerful PHY may not be necessary in order to achieve good performance. Indeed [9] demonstrates that even a low-bit rate, minimal-cost infrared optical modem can substantially improve the performance of underwater acoustic networks.

The variable-depth moored nodes presented in [10] combine acoustic communications under water and radio communications above the water surface by automatically finding a balance between the energy required for the node to surface and the energy consumption of underwater acoustics. The work in [11] employs multi-modal optical and acoustic communications in a clustered optical underwater network. Specifically, the long range characterizing acoustic communications is exploited for cluster formation and management, whereas intra-cluster communications take place through optical connections.

Hybrid acoustic/optical multi-modal networks are considered for the transmission of video streams in [12]. The solution is shown via simulations to outperform optical and acoustic communications alone. In [13], the authors assume that sensor data generated by underwater nodes loses value over time and optimize the trajectory of an AUV to maximize the value of the information that it retrieves from the sensors. Sensor-to-AUV data upload takes place through an optical connection, whereas the sensors notify the AUV of new data using control packets through acoustic connections. Simple context-based PHY switching schemes are considered in [14] to manage multi-modal communications for the control of a remotely operated vehicle (ROV). The authors show how the ROV-controller link can benefit from each technology and how well the ROV reacts to commands. In [15], the free-access DESERT underwater framework [16] is used to implement and evaluate multi-modal communications in a diver cooperation scenario.

The authors in [3] propose hybrid acoustic/optical communications to coordinate swarms of autonomous underwater vehicles and to transfer information among swarm components. A bilingual modem concept was implemented in [17] using a custom re-configurable underwater acoustic modem. Two modulation schemes were employed, namely the NATO standard JANUS and a higher-rate multi-level FSK format.

While the above works were key to introduce and improve the use of multi-modal technologies in underwater networks, the respective approaches typically imply the use of heuristic solutions to divide the data between the PHY technologies. As a result, the multi-modal network is not optimally utilized. Considering this challenge, in this paper we present a routing algorithm that aims to maximize the multi-modal network goodput, under constraints that divide resources fairly among the nodes and avoid the occurrence of bottlenecks.

III. NETWORK MODEL AND OPTIMAL ROUTING

A. Key idea

The key idea behind our distributed routing scheme is that the available multi-modal links should be fully exploited, while at the same time avoiding that some relays become bottlenecks for the routing process. To do so, the nodes should *i*) avoid forwarding an excessive amount of traffic towards the relays upstream and *ii*) favor nodes with fewer valid routes to the sink during the data relaying process. We achieve this by having upstream relays distribute minimal information about their current transmission backlog, so that the nodes downstream can estimate how much data it is safe to forward. This

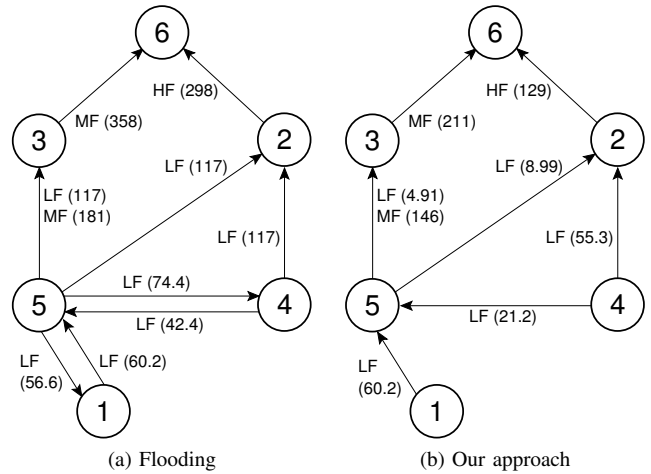


Fig. 1. Key idea: example of per-link transmission rates (kbit/s) to transmit the same number of packets using flooding (a) and our algorithm (b).

allows each network node to separately solve an optimization problem, and to find the number of bits to be transmitted to its neighbors through each multi-modal link. We note that our approach does not resort to flooding, as we explicitly want to avoid unnecessary redundant transmissions of the same data.

Even without topology information, this approach balances traffic much better, e.g., with respect to a baseline algorithm that floods data through all available technologies. This can be observed in Fig. 1 for a 6-node network. The figure shows results for multi-modal routing through flooding and for a version of our algorithm that achieves fairness with only one-hop topology information. Three acoustic PHY technologies are used: a low-frequency, low-rate technology (LF), a faster technology working at intermediate frequencies (MF) and a high-frequency high-rate technology (HF). We ran the two schemes for the transmission of a fixed number of packets and stopped their operation once all packets are received in a scenario configured similarly to the one presented later in Section IV. In parentheses, we show the transmission rates (in kbit/s) needed to achieve the successful transmission of all packets. Compared to flooding, our algorithm requires much less intense transmissions, and thus proves to have much less overhead. Moreover, the results show that the transmission rates of our algorithm are more balanced across the links. In the following, we describe the routing algorithm in detail.

B. Preliminary considerations and notation

We consider a multi-modal converge-casting network, where data traffic is routed towards a sink node. We desire to obtain good performance in all key aspects of the multi-modal network. In particular, we are interested in minimizing the end-to-end transmission delay while maximizing the network goodput, obtaining low packet transmission overhead (to accommodate energy limitations), and high exploitation of the available network links, so as to keep network traffic flow smoothly. Our optimal multi-modal routing (OMR) scheme, presented below, has the goal of achieving a favorable tradeoff between the above quality measures.

We assume that our underwater network is composed of a set \mathcal{N} of multi-modal nodes, where $|\mathcal{N}| = N$, and that the network topology has been already discovered (e.g., see [18]). We however do assume that the process can be subject to errors or inaccuracies due to slow topology changes over time. We will take these errors into account in the design of OMR. Given the outcome of the routing structure discovery, we assume that each node knows the alternatives it has to forward a packet towards the sink D . Accordingly, for each node, we call \mathcal{Y}_i the set of upstream neighbors of i , i.e., \mathcal{Y}_i contains all one-hop neighbors of i that can advance packets one further hop towards D and $\tilde{\mathcal{Y}}_i$ the list of *all* one-hop neighbors of i .

Call \mathcal{T}_i the set of PHY technologies available to node i . A node can communicate using any technology available to it and to the addressed receiver. The list of available technologies may vary over time: let $\mathcal{T}_{i,j}^\tau$ be the set of technologies that i can use to transmit to j at time τ . We assume that $\mathcal{T}_{i,j}^\tau$ is known to the routing protocol (e.g., through signaling by the underlying MAC protocol). Each node maintains a queue with a list of packets to transmit. Denote the bits in node i 's queue at time τ as $P_{i,\tau}$. Define $R_i^\tau(j, t)$ as the number of bits in $P_{i,\tau}$ that will be sent by node i to node $j \in \mathcal{Y}_i$ using technology t at time τ . The objective of the routing algorithm is to find optimal values for $R_i^\tau(j, t)$, under a constraint on the total number of bits that can be transmitted by i using technology t over a user-defined time span u , denoted as $C_{i,t}^{(u)}$. We will indicate these optimal values as $\widehat{R}_i^\tau(j, t)$.

The optimization is to be carried out using the information available at node i or passed on by its upstream neighbors. In particular, we assume that node i knows: \mathcal{Y}_i and $\mathcal{Y}_j \forall j \in \mathcal{Y}_i$; $P_{i,\tau}$ and $P_{j,\tau'} \forall j \in \mathcal{Y}_i$, where $\tau' < \tau$ is a time epoch that refers to a transmission carried out by node j immediately preceding the current epoch τ ; and $C_{j,t}^{(u)} \forall j \in \mathcal{Y}_i$.

C. Routing algorithm

Our goal is to determine $\widehat{R}_i^\tau(j, t)$ and allocate data loads to the different technologies t available to node i . To compute $\widehat{R}_i^\tau(j, t)$, node i solves the following problem for each upstream neighbor in \mathcal{Y}_i (for clarity, let j be one of these upstream nodes):

$$\widehat{R}_i^\tau(j, t) = \arg \max_{R_i^\tau(j, t)} \sum_{j \in \mathcal{Y}_i} \sum_{t \in \mathcal{T}_{i,j}^\tau} R_i^\tau(j, t) \quad (1a)$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{Y}_i} \sum_{t \in \mathcal{T}_{i,j}^\tau} R_i^\tau(j, t) \leq P_{i,\tau}; \quad (1b)$$

$$\sum_{t \in \mathcal{T}_{i,j}^\tau} R_i^\tau(j, t) \leq \Delta_j; \quad (1c)$$

$$R_i^\tau(j, t) \leq C_{i,t}^{(u)} F_j(i). \quad (1d)$$

Constraint (1b) means that the bits transmitted across all technologies shall not exceed the remaining number of bits in the queue at node i . Constraint (1c) takes into account the aggregated capacity of i 's upstream neighbor node j , denoted as Δ_j . Finally, constraint (1d) means that the number of bits transmitted through either technology should be limited by

node i 's fair share $F_j(i)$ of j 's upstream link capacity, where $F_\ell(j) = 0$ if node ℓ has nothing to transmit.

Note that $\sum_i F_j(i)$ can exceed 1. This is because condition (1d) only applies if a node i has more possible relays to the sink than node j : in this case, i should divide its transmissions while considering the relay options of other neighbors of j . Also note that since we limit ourselves to a distributed solution, node i has typically no way to ascertain the technology used over link $j \rightarrow k$, $k \in \mathcal{Y}_j$. Instead, we perform technology allocation only hop-by-hop. As a result, $F_j(i)$ is not related to the used technology t .

The quantities required to evaluate the constraints are fully determined by node i . Node i is assumed to know the capacity of its one-hop links, Δ_j , its available technologies, and its different paths to the sink. However, $F_j(i)$ must still be computed, as will be detailed in the following.

We start with the computation of $F_j(i)$. The upstream transmission resources of node j are assigned to a downstream neighbor i depending on the number of node-disjoint routes towards the destination D available to i , indicated with $L_j^{(i)}$, where the subscript j indicates that $j \in \mathcal{Y}_i$, and that it is being considered as a next hop. The rationale behind the resource assignment strategy is that if a downstream neighbor m of j can reach the destination only via a route that passes through node j , such node m should be given higher priority in the use of j 's upstream transmission resources. Formally, define

$$\widetilde{L}_j^{(i)} = \sum_{\ell \in \mathcal{Y}_j} L_j^{(\ell)} - L_j^{(i)}. \quad (2)$$

If $\widetilde{L}_j^{(i)} = 0$, then we immediately set $F_j(i) = 1$, as j is the only neighbor of i that can relay packets towards D . Otherwise, $F_j(i)$ is computed as $F_j(i) = \widetilde{L}_j^{(i)} / (\sum_{\ell \in \mathcal{Y}_j} \widetilde{L}_j^{(\ell)})$, where it is understood that $i \notin \mathcal{Y}_j$, i.e., i is not an upstream neighbor of j . Note that this is a way to "fairly" allot more resources to nodes with fewer available routes, not a means to split the capacity of node j 's links towards its upstream neighbors, which is instead done distributedly via (1c). Instead, we allow nodes with a single forwarding opportunity to convey all traffic there, while nodes with additional opportunities should split their traffic through all available routes.

To compute $L_j^{(i)}$ in (2), we distinguish two cases: a) full topology-informed fair share computation, in case perfect topology information is available to i (dubbed OMR-FF); and b) one-hop topology-informed fair share computation, otherwise (OMR-PF, which is also the version employed for the results in Fig. 1b). In the case of OMR-FF (case a), node i is aware of the full network graph, and can thus compute the number of disjoint routes available to itself and its neighbors, thereby exactly deriving its own fair share of resources. In the case of OMR-PF (case b), only one-hop topology information is available to i . In this case, $L_j^{(i)}$ is estimated as

$$L_j^{(i)} = |\mathcal{Y}_j| - \sum_{w \in \tilde{\mathcal{Y}}_i \cup \mathcal{Y}_j} \mathbb{1}[\mathcal{Y}_w = \{i, j\} \vee \mathcal{Y}_w = \{i\} \vee \mathcal{Y}_w = \{j\}] \quad (3)$$

where $\mathbb{1}[p]$ evaluates to 1 if predicate p is true. Eq. (3) means that, as a best effort, $L_j^{(i)}$ is assumed to be equal to the number

of downstream neighbors of k , decreased by one for each node w that has only i , j , or both as upstream neighbors, which may occur due to erroneous topology information. This avoids that node i gets an excessive resource share. We note that the computation of $L_j^{(i)}$ in (3) is not carried out if the destination $D \in \mathcal{Y}_j$. In this case, the traffic is always directed to D .

The nodes decide on the number of bits to be transmitted over each link regardless of what packet such bits belong to, and manage their queue in a FIFO fashion. Each transmitted datagram includes the information required by the receiver to properly re-assemble the bits.

IV. SIMULATIONS

We now study the performance of OMR–FF and OMR–PF through numerical simulations. We recall that the difference between the two flavors of OMR affects the way resources are allocated to different nodes over a multi-hop path: namely, OMR–PF tends to be more conservative in terms of link capacity usage than OMR–FF. Except for this aspect, the two OMR flavors behave in the same way.

In the absence of a benchmark routing scheme for multi-modal networks, we compare the performance of the two versions of OMR with that of a *flooding* mechanism, where a node broadcasts all incoming packets through all available technologies. To prevent loops in the flooding scheme, each packet includes a list of visited nodes. A receiver will then avoid broadcasting a received packet if this list includes all its one-hop neighbors. In the flooding method, packets are fragmented according to the maximum datagram length allowed by the technology through which the packet is sent.

A. Quality metrics

We measure the performance in terms of end-to-end transmission delay, per-node goodput, overhead, and per-link transmission rate. Once all fragments of a packet i of node n have been successfully received by the sink, we measure the average end-to-end transmission delay (over all packets and nodes) as

$$\rho_d = \frac{1}{N-1} \sum_{n=1}^{N-1} \frac{1}{R_n} \sum_{i=1}^{R_n} (T_{n,i}^r - T_{n,i}^s), \quad (4)$$

where $T_{n,i}^r$ is the time the full message was received, $T_{n,i}^s$ is the time the message reached the network layer for routing, and R_n is the number of messages sent by node n and received in full by the sink node. For a network run time T_{net} , the per-node goodput is defined as

$$\rho_g = \frac{1}{N-1} \sum_{n=1}^{N-1} \sum_{i=1}^{I_n} \frac{M_{n,i}^r}{T_{\text{net}}}, \quad (5)$$

where $M_{n,i}^r$ is the number of bytes received by the sink for a message i originated by node n , and I_n is the number of messages originated by node n .

Note that $M_{n,i}^r$ from (5) can exceed the number of bytes transmitted by node n , denoted by $M_{n,i}^s$. This case happens when message i or parts of it are sent through several links such that the sink may receive multiple copies. We consider these cases a waste of resource, and refer to this excess of

TABLE I
CHARACTERISTICS OF THE SIMULATED COMMUNICATION TECHNOLOGIES

Technology	Bit rate [bps]	Max range [m]
Low-rate acoustics	1000	3000
Mid-rate acoustics	32000	300
High-rate acoustics	64000	100

copies as a transmission overhead. This overhead is measured by

$$\rho_o = \frac{1}{N-1} \sum_{n=1}^{N-1} \sum_{i=1}^{I_n} U\left(\frac{M_{n,i}^r}{M_{n,i}^s} - 1\right), \quad (6)$$

where $U(x)$ is a step function whose value equals 1 if $x > 0$, and zero otherwise.

Finally, the transmission rate of the link from node n to node m using communication technology t is defined as the ratio between the total number of bytes successfully transmitted through the link, $R_{n,m}^t$, and the run time. Formally, the average per-link transmission rate is

$$\rho_u = \frac{1}{N^t} \sum_{n \in \mathcal{N}^t} \frac{1}{D_n^t} \sum_{m \in \mathcal{D}_n^t} \frac{R_{n,m}^t}{T_{\text{net}}}, \quad (7)$$

where \mathcal{N}^t is the set of the nodes who hold communication technology t , and \mathcal{D}_n^t is the set of the nodes that share a communication link with node n via technology t . Moreover, $N^t = |\mathcal{N}^t|$ and $D_n^t = |\mathcal{D}_n^t|$.

As anticipated in Section III, we wish to minimize ρ_d , and to maximize ρ_g . For energy conservation, we are also interested in minimizing ρ_o . Finally, for better fairness and to minimize the risk of congestion, we are interested in a low ρ_u as well.

B. Simulation setup

Our simulation setup is based on a Monte-Carlo set of 1000 network topologies. In each run, $N = 10$ nodes are placed uniformly at random over an area of 500×500 m² with depth 100 m. Four randomly placed obstacles of length drawn uniformly at random in $[10, 50]$ m may obstruct the line-of-sight between the nodes. Node 10 is the sink. Each of the other nine nodes is equipped with one or more of three communication technologies among low-, mid-, and high-frequency acoustic communications, whose features are based on the 18–34 kHz, the 48–78 kHz, and the 120–200 kHz EvoLogics modems [19]. A summary of these features is provided in Table I, where communication ranges have been conservatively set. Each technology is equipped with probability 0.5, and a node has at least low-frequency acoustics available.

We run each simulation for $T_{\text{net}} = 600$ s. Each of the nine nodes generates traffic at the beginning, according to a Poisson process of rate $\lambda = 3$ packets/min/node. The size of each packet is drawn uniformly at random in the interval $(0, 64]$ kbit. For each packet, the node solves the routing allocation problem presented in Section III. The packet is then segmented according to the solution of the routing problem and sent over the different links according to the determined routing allocation. The various packet segments belonging to the same message are finally reassembled by the sink.




Fig. 2. Simulation results for ρ_d (Eq. (4)) and ρ_g (Eq. (5)). OMR performs better than flooding, with OMR–FF typically achieving the best results.

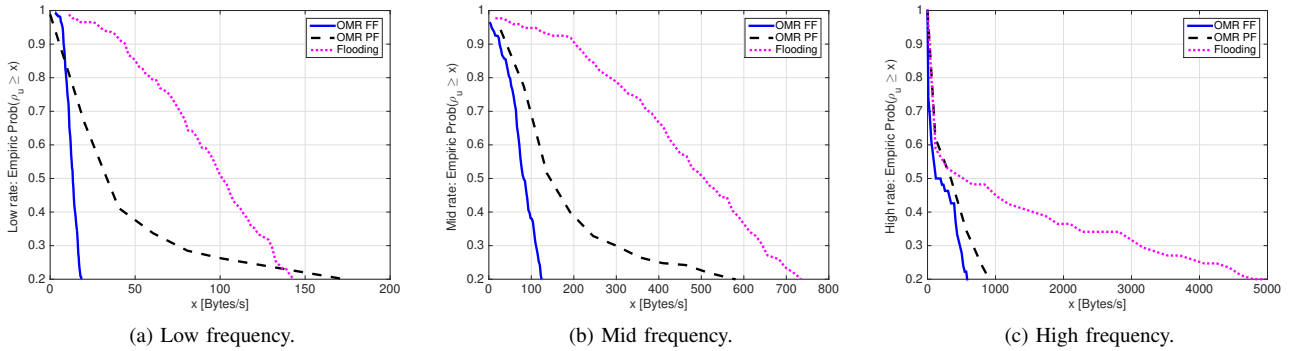


Fig. 3. C-CDF of ρ_u (Eq. (7)). Flooding employs links quite intensively, whereas OMR requires a much lower transmission rate for achieving a similar goodput as seen in Fig. 2b. Moreover, OMR–FF’s performance is more consistent than OMR–PF’s across different topologies.

We consider a binary phase-shift-keying modulation, and a scheduling protocol whereby a node holding a packet transmits it as soon as all its communication technologies are free. Once a packet is received, an acknowledgment is transmitted. To form the full topology information required by OMR–FF and the one-hop link information required by OMR–PF, we refer to the communication ranges in Table I. To calculate the route on the way to the sink (i.e., the sets $\mathcal{Y}_i, \forall i$), we carry out a preliminary route discovery phase, where the sink propagates a discovery packet through the network [18]. The discovered routes are maintained throughout each simulation run.

While in the OMR scheme the full topology (OMR–FF) or the one-hop links (OMR–PF) are assumed to be known, real underwater links are not fully reliable. To simulate this, we calculate the packet error rate (PER) for each link based on the simulated signal-to-noise ratio (SNR) and on the size of transmitted segments. The SNR of the acoustic links is calculated through Bellhop [20, Ch. 3] for shallow waters of depth 100 m, flat sand bottom, fixed sound speed of 1500 m/s, and considering a source level of 170 dB re (1 μ Pa at 1 m). The ambient noise level is set to 40, 30 and 10 dB re (1 μ Pa²/Hz) for the low-frequency, mid-frequency, and high-frequency acoustics, respectively. All transmissions are acknowledged. If a transmission fails, the respective packet is shifted to the end of the message queue for retransmission.

We consider a MAC protocol in which packets are transmitted immediately upon arriving to the MAC layer, unless another transmission or reception is already taking place, and the reception of packets and acknowledgments is determined based on the link SNR and only when no collision occurs with another packet or acknowledgment.

C. Simulation results

Fig. 2a shows the cumulative distribution function (CDF) of the delay ρ_d , which spans more than 200 s. This is because low-capacity links require packets to be segmented into small fragments, which increases the backlog of the nodes. We also observe that flooding’s delay is much higher than OMR’s, and that OMR–FF achieves the lowest delay. The former is due to heavy interference, while the latter is because the availability of topology information makes it possible for the nodes to allocate transmission resources more effectively. However, even without full topology information, the delay of OMR–PF is almost as good as that of OMR–FF.

Next, in Fig. 2b, we show the complementary CDF (C-CDF) of the goodput ρ_g . The figure confirms that OMR manages transmissions effectively, as it achieves about the same goodput as flooding with significantly lower delay (see Fig. 2a). The main reason is again the large number of packet collisions caused by the many transmissions of flooding.

We now consider a comparison of the network fairness, which we measure via the C-CDF of the per-link transmission rate, ρ_u , for low-, mid-, and high-frequency transmissions (Figs. 3a, 3b, and 3c, respectively). We observe that since a node transmits simultaneously over all available links when using flooding, the per-link transmission rate is much higher than OMR's, albeit the goodput is very similar (see Fig. 2b). This is a clear indication that OMR distributes resources more effectively and achieves the same information transfer with fewer transmissions. Fig. 3 also shows that the per-link transmission rate in OMR–FF has a low statistical dispersion, meaning that this metric is maintained stable over different topologies. This is mainly due to full topology information. The lack of this knowledge makes OMR–PF comparatively less effective (its C-CDF for ρ_u has a longer tail) and more prone to bottlenecks (as seen from the larger delay in Fig. 2a).

To comment on the energy efficiency of the three methods, in Fig. 4 we show the overhead ρ_o from (6). While multiple redundant copies of all messages are received with the flooding scheme, in the two OMR versions the sink receives extra copies only for about 8% of the messages. Comparing the goodput performance from Fig. 2b and the overhead results of OMR–FF and of OMR–PF, we observe almost identical performance, and that only a few redundant transmissions exist. This result works in slight favor of OMR–PF, which requires less knowledge and is thus more distributed.

V. CONCLUSIONS

In this paper, we considered the network operation of multi-modal systems, a technology which provides great benefit to underwater networks. We proposed OMR that, to the best of our knowledge, is the first optimal routing protocol to be specifically designed for multi-modal underwater networks. Our protocol leverages either full or local topology knowledge to decide how to distribute traffic over links using the available communication technologies. This is achieved in a way that does not congest the relays upstream, and reserves more link resources for the nodes with fewer routing opportunities. We analyzed the performance of OMR by means of simulations. Our results show that OMR leverages the available technologies (and, if available, topology information) to deliver data reliably without congesting the network.

ACKNOWLEDGMENT

This work has been supported in part by the US Office of Naval Research under Grant no. N62909-14-1-N127.

REFERENCES

- [1] M. Chitre, S. Shahabudeen, and M. Stojanovic, "Underwater acoustic communications and networking: Recent advances and future challenges," *Marine Tech. Soc. Journal*, vol. 42, no. 1, pp. 103–116, 2008.
- [2] N. Farr, A. Bowen, J. Ware, C. Pontbriand, and M. Tivey, "An integrated, underwater optical/acoustic communications system," in *Proc. MTS/IEEE OCEANS*, Sydney, Australia, May 2010.
- [3] C. Moriconi, G. Cupertino, S. Betti, and M. Tabacchiera, "Hybrid acoustic/optic communications in underwater swarms," in *Proc. MTS/IEEE OCEANS*, Genova, Italy, May 2015.
- [4] X. Che, I. Wells, G. Dickers, P. Kear, and X. Gong, "Re-evaluation of RF electromagnetic communication in underwater sensor networks," *IEEE Commun. Mag.*, vol. 48, no. 12, pp. 143–151, Dec. 2010.

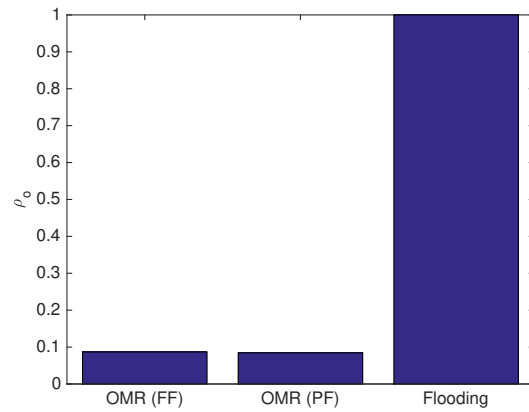


Fig. 4. Average overhead ratio ρ_o from (6). OMR is more efficient than flooding. OMR–FF and OMR–PF perform similarly.

- [5] "Wireless For Subsea Seatooth," accessed: Feb. 2016. [Online]. Available: <http://www.wfs-tech.com/index.php/products/seatooth/>
- [6] F. Campagnaro, F. Favaro, P. Casari, and M. Zorzi, "On the feasibility of fully wireless remote control for underwater vehicles," in *Proc. Asilomar Conf. SS&C*, Pacific Grove, CA, Nov. 2014.
- [7] H. Dol, M. Colin, P. van Walree, and R. Otnes, "How smartphone industry made UComms easier: Moore's law goes underwater," in *Proc. MTS/IEEE OCEANS*, Genova, Italy, May 2015.
- [8] I. Vasilescu, K. Kotay, D. Rus, P. Corke, and M. Dunbabin, "Data collection, storage and retrieval with an underwater optical and acoustical sensor network," in *Proc. ACM Sensys*, San Diego, CA, Nov. 2005.
- [9] A. Tennenbaum, M. Dyakiw, J. H. Cui, and Z. Peng, "Application of low cost optical communication systems to underwater acoustic networks," in *Proc. IEEE MASS*, Philadelphia, PA, Oct. 2014.
- [10] M. O'Rourke, E. Basha, and C. Detweiler, "Multi-modal communications in underwater sensor networks using depth adjustment," in *Proc. ACM WUWNet*, Los Angeles, CA, Nov. 2012.
- [11] R. Kastner, A. Lin, C. Schurgers, J. Jaffe, P. Franks, and B. S. Stewart, "MURAO: A multi-level routing protocol for acoustic-optical hybrid underwater wireless sensor networks," in *Proc. IEEE SECON*, Seoul, South Korea, Jun. 2012.
- [12] S. Han, R. Chen, Y. Noh, and M. Gerla, "Real-time video streaming from mobile underwater sensors," in *Proc. ACM WUWNet*, Rome, Italy, Nov. 2014.
- [13] S. Basagni, L. Bölöni, P. Gjanci, C. Petrioli, C. A. Phillips, and D. Turgut, "Maximizing the value of sensed information in underwater wireless sensor networks via an autonomous underwater vehicle," in *Proc. IEEE INFOCOM*, Toronto, Canada, Apr. 2014, pp. 988–996.
- [14] F. Campagnaro, F. Guerra, F. Favaro, V. S. Calzado, P. Forero, M. Zorzi, and P. Casari, "Simulation of a multimodal wireless remote control system for underwater vehicles," in *Proc. ACM WUWNet*, Washington, DC, Oct. 2015.
- [15] F. Campagnaro, F. Guerra, R. Diamant, P. Casari, and M. Zorzi, "Implementation of a multimodal acoustic-optic underwater network protocol stack," in *Proc. MTS/IEEE OCEANS*, Shanghai, China, Apr. 2016.
- [16] P. Casari *et al.*, "Open-source suites for underwater networking: WOSS and DESERT Underwater," *IEEE Network, special issue on "Open Source for Networking: Development and Experimentation"*, vol. 28, no. 5, pp. 38–46, Sep. 2014.
- [17] R. Petrocchia, G. Cario, M. Lupia, V. Djapic, and C. Petrioli, "First in-field experiments with a bilingual underwater acoustic modem supporting the JANUS standard," in *Proc. MTS/IEEE OCEANS*, Genova, Italy, May 2015.
- [18] R. Diamant, R. Francescon, and M. Zorzi, "Efficient link discovery for underwater networks," in *Proc. UCOMMS*, Lercici, Italy, Aug. 2016.
- [19] "EvoLogics Underwater Acoustic Modems," last time accessed: Oct. 2016. [Online]. Available: <https://www.evologics.de/en/products/acoustics/index.html>
- [20] F. Jensen, W. Kuperman, M. Porter, and H. Schmidt, *Computational Ocean Acoustics*, 2nd ed. New York: Springer-Verlag, 1984, 2nd printing 2000.