

A Node Discovery Protocol for Ad Hoc Underwater Acoustic Networks

Ashish Patil and Milica Stojanovic

Dept. of Elec. and Computer Eng., Northeastern University, Boston, MA 02115

Abstract—A neighbor discovery protocol is proposed for initializing an ad-hoc deployable autonomous underwater system in which the nodes (floating sensors, crawlers, vehicles) have no knowledge of the system topology upon deployment. The goal is to establish multi-hop, minimum-power acoustic communication links over a given coverage area. The protocol uses distributed power control and random access to provide connectivity within a finite power budget and without global synchronization. Physical laws of acoustic propagation are taken into account, namely the distance- and frequency-dependent path loss, as well as large-scale fading which is modeled via log-normal distribution. Simulation results quantify the energy consumption, time to completion, and system reliability in the presence of fading. The key features of the protocol are simplicity of implementation and efficient use of power.

I. INTRODUCTION

Ad hoc deployable systems with both mobile and slowly drifting nodes are envisioned for applications such as data gathering and instrument maintenance, search and survey missions for both military and commercial applications, and exploratory tasks that serve basic sciences. The key technologies that will make such applications possible – vehicles, sensors and communications – are today mature enough to warrant their integration into fully operational systems with a high degree of autonomy.

The fundamental aspects of underwater acoustic networking have been highlighted in several recent publications, e.g. [1], [2], [3]. In the past years there have been major developments on the medium access control (MAC) layer [4], [5], [6], as well as on the routing layer [7], [8], [9]. Many of the proposed protocols focus on the design that is not ignorant of the high acoustic latency, but strives to overcome it in an efficient manner or even take advantage of it. Cross-layer design, e.g. in the form of coupling the power control with MAC and routing [9], also plays an important role in acoustic systems.

Unlike in a fixed network, the nodes in an ad hoc network are deployed in a (more or less) random fashion, and have no knowledge of their neighbors' IDs or locations upon deployment – an assumption on which an efficient acoustic MAC and routing protocol suite typically relies. The network's first task is thus to establish the communication links. The information gathered about the network topology during initialization (the neighbors' IDs and locations) can be used to build routing tables, which can later be dynamically updated.

Node discovery has been extensively investigated for radio networks (e.g. [10]- [14]) but work in the field of underwater acoustic networks remains scarce. A specific neighbor discov-

ery procedure for an underwater acoustic network was first discussed in [15]. This reference proposed a protocol based on polling by a master node in a centralized configuration, in which the channel access is regulated through code division-multiple access. Ref. [16] proposed a similar neighbor discovery procedure. In Ref. [17] an experimental deployment is described in which initial neighbor discovery is performed via a master node.

In this paper, we propose a decentralized node discovery procedure which is suitable for large deployment areas that cannot be spanned in a single hop within the constraints of finite power. The nodes operate in a distributed manner, i.e. without a central station or a-priori regulation of the channel access. We consider a random access environment in which there is no a-priori division of the available resources, such as time, frequency, or code-division multiple access, and no scheduling. Random access is advantageous from the viewpoint of simplicity of implementation which requires no global synchronization for scheduling.

The paper is organized as follows. In Sec.II we briefly summarize the acoustic propagation model used for system design and its performance evaluation. In Sec.III, we present the node discovery protocol, including a detailed illustration of the protocol operation. Performance analysis is presented in Sec.IV, which contains simulation results for both the nominal case with no fading and for the realistic case of fading, as well as comparisons with the benchmark case of broadcast initialization. Finally, conclusions are summarized in Sec.V.

II. ACOUSTIC PROPAGATION MODEL

The received signal strength in an acoustic channel is determined by the channel geometry and propagation conditions, as well as by the frequency occupancy of the signal. If we denote by $H(d, f)$ the channel transfer function between a transmitter and receiver separated by a distance d , the average signal power, received in response to a transmit power P_T , allocated uniformly across a signal bandwidth B centered at frequency f_c , is given by

$$P_R = \frac{P_T}{B} \int_{f_c-B/2}^{f_c+B/2} E\{|H^2(d, f)|\} df \equiv P_T \cdot G(d) \quad (1)$$

where the expectation is taken over the fast, small-scale channel variations, but not over the slower, large-scale variations. The factor $G(d)$ thus represents the large-scale channel gain (inverse of attenuation). Because the propagation conditions vary in time and with small displacements of the transmitter

and receiver around their nominal locations, the channel gain $G(d)$ varies as well. These variations are usually perceived as random, and we model them via a log-normal distribution (see [18] and references therein). Specifically, we model the gain, expressed in dB as $g(d) = 10\log G(d)$, as

$$g(d) = \bar{g}(d) + x \quad (2)$$

where $\bar{g}(d)$ is the nominal gain, obtained for a fixed link geometry, and x is a normally distributed random variable with mean zero and variance σ^2 . The fading strength is determined by the ratio $|\bar{g}(d)|/\sigma$, which we will later refer to as the fading parameter A/σ .

The nominal gain depends on the particular multipath composition of the channel, which is in turn determined by the deployment location. In order to maintain generality, in what follows we will assume an ideal case in which there is no multipath, i.e. we will take

$$E\{|H^2(d, f)|\} = \frac{1}{A(d, f)} \quad (3)$$

where $A(d, f)$ is the nominal (basic) acoustic propagation loss observed over distance d at frequency f . Extensions to particular multipath settings are straightforward.

The basic loss is modeled as [19]

$$A(d, f) = A_0 d^k a(f)^d \quad (4)$$

where A_0 is a scaling constant, k is the spreading factor (1 for cylindrical, 2 for spherical spreading geometry), and $a(f)$ is the absorption coefficient.

The noise power, observed in the frequency range of interest, is given by

$$P_N = \int_{f_c - B/2}^{f_c + B/2} N(f) df \quad (5)$$

where $N(f)$ is the power spectral density (p.s.d.) of the background noise. We approximate $N(f)$ as [20]

$$N(f) \approx N_0 f^{-\eta} \quad (6)$$

where N_0 is a constant that can be measured for a particular system location, and the factor η models the p.s.d. decay, which usually occurs at 18 dB/decade.

The key parameter for link design is the signal-to-noise ratio (SNR). Specifically, we assume that successful packet detection is possible if the SNR is above a pre-specified threshold,

$$\text{SNR} = \frac{P_R}{P_N} \geq \text{SNR}_0 \quad (7)$$

The SNR condition implies that the transmit power needed to close a link of distance d is

$$P_T \geq \text{SNR}_0 \frac{P_N}{G(d)} \quad (8)$$

The presence of random component in the channel gain $G(d)$ makes it impossible to adjust the transmit power exactly. For this reason, we define the nominal transmit power needed to

close a link of distance d as

$$P(d) = \text{SNR}_0 \frac{P_N}{G(d)} \quad (9)$$

We will use this value later when we discuss power control.

III. THE DISCOVERY PROTOCOL

Our goal is to design a protocol that is efficient in terms of energy consumption as well as the time it takes to complete. The protocol must establish full connectivity, i.e. it must guarantee, within the power constraint and in the absence of fading, a path through the network between any two participating nodes. (A node that cannot be reached by another node at the maximal power level is considered not to be participating.)

The following assumptions are made:

- There is a finite number of nodes, N .
- The nodes know the number N .
- A nodes knows its location with a certain accuracy.
- Nodes can vary their transmission power in discrete levels between some P_{min} and P_{max} .
- Nodes operate in half-duplex fashion.
- Nodes have enough processing power and memory to perform simple calculations and store information about the network topology.
- Nodes have notion of relative time with certain accuracy, but do not need to be synchronized to a global clock.

The protocol is designed assuming a quasi-stationary scenario in which the network topology does not change significantly during the discovery. This assumption does not imply that the nodes may not move, only that the local neighborhood topology does not change faster than the nodes can learn it. Given a typical speed of an autonomous underwater vehicle (AUV) of a few meters per second, and a temporary confinement area of several hundred square meters, this is a reasonable assumption.

The protocol proceeds in cycles, each led by a single node. The leader broadcasts a message and waits to receive replies. Those nodes that hear the leader, reply using a simple Aloha mechanism. Random access may also be favored for later, regular network operation [21], in which case the transition from discovery will be seamless. At the end of a cycle, leadership is passed to another node, or retained if an increased power level is required to find a new neighbor. Power control is implemented in discrete levels as described in [9], and accounts for both the distance-dependent transmission loss and the channel fading. As the leadership is forwarded, a sign-up sheet is circulated among the leaders. The procedure ends when all the nodes have signed up, or when all the nodes have exhausted their maximum power level. Once the discovery phase is over, normal network operation can begin.

Note also that once the regular network operation has begun, the discovery protocol may continue to run in the background, perhaps at a lower pace. By doing so, a *recovery* procedure can be put in place to guard against the loss of nodes. Such situations can occur in a mobile setting, if a node approaches a

zone of poor coverage. Its connection to the current neighbors will then be severed, but so long as the network is aware of the loss and its nodes have the ability to move, it can initiate recovery, similarly as it initiated the original discovery.

A. Network Topology and Power Control

The area over which the network is deployed is assumed to be a square of side D . A grid is imposed onto this area so as to divide it into smaller cells as shown in Fig.1. Accordingly, we assume that a node is placed randomly within its cell.




Fig. 1. Nodes are located randomly within a grid.

Every node has a finite number of power levels L that correspond to distances

$$d_l = d_0 + l\Delta d, l = 0, \dots, L-1 \quad (10)$$

Given a total of N nodes, the minimum and the maximum distance are set to (see Fig.1)

$$d_0 = d_{min} = \frac{D}{\sqrt{N}} = \frac{1}{\sqrt{\rho}} \quad (11)$$

and

$$d_{L-1} = d_{max} = d_{min} \sqrt{5} \quad (12)$$

where ρ denotes the node density, and we assume without loss of generality that N is a square number, i.e. $N=4, 9, 16, 25, 36$, etc. Given a desired number of power levels, the step Δd is determined as

$$\Delta d = \frac{d_{max} - d_{min}}{L-1} \quad (13)$$

Transmission power required to provide a target SNR_0 at a particular distance d_l is now obtained from the expression (9). We denote the l -th power level by $P_l = P(d_l)$. Conversion from acoustical power expressed in dB re μPa^2 to electrical power expressed in dB re W is performed by subtracting 172 dB [22].

A node that has no prior knowledge of the power needed to reach an intended receiver will always start from the lowest power level P_0 . The power is increased only if necessary, as dictated by the discovery process. This design targets efficient use of power in a network with large coverage where it would be impossible or wasteful of power for a single node to reach out to all the other nodes.

B. Signaling Format and the Discovery Cycle

The discovery packet contains the following information: [type: D, source, destination, power level]. Since the discovery procedure begins with node 1 transmitting at the lowest power level, the first discovery packet to be sent is of the form [type: D, source: $N1$, destination: any node, power level: P_0].

After transmitting the discovery packet, the leader switches to the listening mode, waiting for replies. If there are no replies, the leader increases its power and repeats the discovery packet. The waiting time is set in accordance with the propagation delay corresponding to the targeted distance. Specifically, the amount of time it takes to complete a transaction at power level l is

$$T_l = (T_P + \tau_l) + (T_R + \tau_l) + T_G \quad (14)$$

where T_D is the duration of the discovery packet, T_R is the duration of the reply packet, $\tau_l = d_l/c$ is the propagation delay over distance d_l (c is the speed of sound), and T_G is the guard interval introduced to account for any errors. Namely, because the time-varying propagation conditions can cause the SNR to deviate from the predicted value, a signal transmitted at power P_l may not be detected by a node closer than d_l , or it might be detected by a node farther than d_l . The latter case will cause the reply to arrive later than expected. To guard against this anomaly, a guard time is introduced.

Nodes that hear the discovery packet at sufficient SNR read from it the leader's information (ID, power level) and reply. The reply packet is of the form [type: R, source, destination: leader, power level, current location].

After acquiring the information about its nearest neighbors, the leader stores this information in a list of contacts that it will keep for future use. It then puts its ID on the sign-up sheet and passes the leadership on in a packet called the end-of-cycle packet.

The outgoing leader also selects the next leader from its list of contacts, and includes this information into the end-of-cycle packet. When there is more than one node in the leader's list of contacts, the new leader is chosen as the closest node that has not yet been a leader. The reason behind this strategy is that shorter links require less power (although they may prolong the discovery).

C. Channel Access

Channel access is regulated using an Aloha-style protocol with carrier sensing. Upon receiving a discovery packet, nodes reply to the leader immediately if they sense the channel free. Because the receiving nodes may not be aware of each other's existence, their replies may collide at the leader (there are no collisions among the discovery packets or the end-of-cycle packets). Note that collisions will be possible only if the replying nodes are approximately at the same distance from the leader. For example, with 24 bit reply packets and a bit rate of 10 kbps, there will be no collisions if the distances of replying nodes from the leader differ by more than 1.8 m. Hence, the chances that reply packets collide at the leader are small.

Nonetheless, the protocol has to account for such a possibility, and it does so through a collision recovery procedure.

1) *Collision recovery*: If two (or more) reply packets collide, the leader remains unaware of the identity of the nodes involved in the collision, although it is aware of the collision itself. To learn the identity of the colliding nodes, the leader creates a new packet called the collision recovery packet, which contains the following information [type: CR, source, destination: any node, power level, list-of-contacts]. The leader then sends this packet at a power level equal to that of the last transmitted discovery packet. The receiving nodes first check the leader's list-of-contacts. If a node finds itself on this list, it will simply ignore the CR packet since it will know that the leader has correctly received its reply. The remaining nodes whose names do not appear in the CR packet will conclude that they have been involved in a collision, and that their replies have to be re-transmitted. This time, however, the nodes will not transmit immediately. Instead, each node will wait for a random back-off time and re-transmit only then. As in any back-off procedure, the additional random delay will help to reduce future collisions.

2) *End-of-cycle*: For the discovery protocol to proceed correctly it is necessary that the end-of-cycle packet (ECP) be received by the next leader. Because of fading, however, the received signal strength deviates from the design value, and the ECP can be lost. Fig.2 illustrates such a situation.

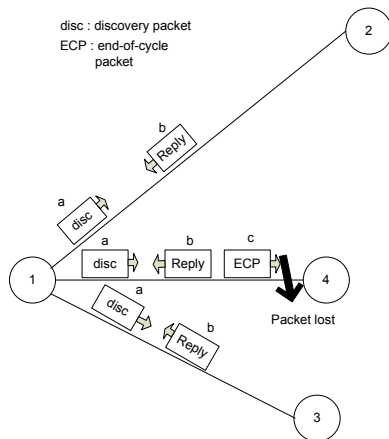


Fig. 2. End-of-cycle packet (ECP) can be lost due to fading, causing the discovery procedure to stop. A possible remedy for this situation is to transmit the ECP multiple times (if one transmission experiences fading, chances are that another will not).

Loss of the ECP will cause the discovery procedure to stop, leaving some nodes undiscovered. The effect of fading will thus be to degrade the system performance. While correct reception can never be *guaranteed* under random propagation conditions with finite power budget, the protocol can be adjusted to increase the *reliability* with which the end-of-cycle packet is received. In particular, the protocol can mandate that the ECP be acknowledged, or that it be repeated a sufficient number of times. Since acknowledgments can also get lost (leading to further undesired effects) we will focus here on

the ECP repetition as an example design. In this design, the ECP is repeated for a pre-specified number of times. The rate of repetitions should be chosen in accordance with the speed at which the channel conditions change (coherence time of the fading process). While many different protocol adjustments are possible to guard against fading, the simple ECP repetition suffices to illustrate the principles.

D. Performance Illustration

Fig.3 shows an example of a newly deployed network which lacks structure, whereas Fig.4 shows the network after completion of discovery. All the nodes in Fig.4 know their neighbors and the power level required to reach them.




Fig. 3. A newly deployed network has no structure, i.e. there are no connections between the nodes.




Fig. 4. Node connections are established during discovery.

Referring to Figs.3 and 4, we now illustrate the steps of the discovery procedure. For simplicity, we restrict our attention to the ideal case with no fading. In this case, the discovery protocol has to guarantee that all the nodes that are within each other's reach are accounted for by the end of the discovery phase. In doing so, the goal is to limit the total power consumption, i.e. not all the connections have to be discovered, but full connectivity has to be provided.

We set the total number of power levels to be $L = 4$ (P_3 is the maximum power that a node can have). With this setting, it takes 17 cycles to complete the discovery of $N=9$ nodes in our example. Table I provide a detailed summary of the system state at the end of each cycle.

In this example, during cycle 1, node 1 does not find any neighbors at power P_0 , but in cycle 2 it finds nodes 2 and 4 at P_1 . It chooses node 2 to be the leader of the next cycle. The state of the system at this point is specified by the second row of Table I, which lists the cycle number (2), the current leader and its power level (node 1 at P_1), the current leader's list of contacts and the power levels needed to reach them (node 2 at P_1 and node 4 at P_1), the next leader (node 2), the sign-up sheet that is being passed on (it contains only node 1 at this point), and the return address (for the moment, this field is empty).

The return address is the address of the node that originally elected the current leader. This address is to be used when the current leader has reached the maximum power, but cannot find any new leaders (any nodes who have not been leaders before). In such a situation, the current leader puts an end mark next to its name (indicated by the superscript "e" in the table) and returns the leadership. Note that it is the *original*, not the last elector, to whom the leadership should be returned. The end mark informs the other nodes that this node is a dead-end, i.e. that there is no point in asking it again to be the leader, because it has exhausted its powers. In our example, such a situation occurs in cycle 11. During the previous cycle, node 4 (the leader) has reached P_2 , but found no candidates for a new leader. It thus increases the power to the maximal level P_3 , but does not discover any new neighbors. Node 4 then signs up with an end mark, and returns the leadership to node 7 (whom it has marked as its return address in cycle 8, when it became the leader for the first time).

The procedure ends when the sign-up sheet is full, or when all the nodes have exhausted their capabilities. The latter case may occur if there are nodes that are outside of the maximal power reach of any other node. These nodes do not belong to the network in the sense in which we have defined it, i.e. within the given power constraint.

Once the discovery procedure has ended, all the nodes in the network can begin regular operation. At this time, not all the nodes may be aware of the fact that the discovery phase is over. In order to inform the other nodes of the completion of discovery, the last leader may commence regular network operation by sending an end-of-discovery packet to its return address, who will then propagate it downstream. This communication will occur in the form of regular data packets, and will be carried through the network using the channel access method of choice. The end-of-discovery packet may contain the sign-up sheet, or simply the statement that the discovery phase is complete. The exact contents of this packet should be determined so as to best serve the system requirements (localization, routing).

IV. PERFORMANCE ANALYSIS

The protocol performance is assessed via simulation, conducted using AUVNetSim, a discrete event simulator written in Python [23]. The following system parameters were used:

- square area of side $D = 10$ km
- center frequency $f_c = 10$ kHz
- bandwidth $B = 10$ kHz, and bit rate 10 kilobits/sec
- target $\text{SNR}_0 = 25$ dB
- background noise level $N_0 = 50$ dB re $\mu\text{Pa}^2/\text{Hz}^1$
- number of power levels $L = 4$
- number of nodes $N=4, 9, 16, 25, 36, 49, 64, 81$
- fading parameter $A/\sigma=17, 20, 24, 30$
- reply packet size: 24 bits
- number of ECP transmissions: 1 without fading; 4, 5 or 7 with fading.

The results are averaged over 1000 random deployments of nodes (a node is placed within its cell uniformly at random, independently of other nodes). The performance is measured in terms of the energy and time consumed during the discovery. Below, we illustrate the results obtained for the nominal case without fading, as well as those obtained in the presence of fading. Fading is assumed to occur independently on different links and at different times of transmission.

A. Ideal Case (No Fading)

Fig.5 shows the average total energy consumption versus the discovery time, obtained for a varying number of nodes in the nominal case without fading. The total energy is the sum of energy spent on transmission, reception, and idle listening. While the transmission energy depends on the inter-node distances, the energy spent on receiving does not. The receiving power is set to 0.1 W for each node, which amounts to about 10, 20, 30 and 50 J for 36, 49, 64 and 81 nodes, respectively.




Fig. 5. Total energy (average) vs. the time it takes to complete the discovery.

¹Noise level depends on the particular system location and can vary by several tens of dBs. Considering a higher noise level will not alter the general performance *trends*, only the absolute values.

Cycle	Leader	Leader's contacts	Sign-up sheet	Next	return
1	1 (P_0)	-	1	1	-
2	1 (P_1)	2(P_1), 4(P_1)	1	2	-
3	2 (P_0)	5(P_0)	1,2	5	1
4	5 (P_0)	2(P_0), 8(P_0)	1,2,5	8	2
5	8 (P_0)	7(P_0), 5 (P_0)	1,2,5,8	7	5
6	7 (P_0)	8(P_0)	1,2,5,8,7	7	8
7	7 (P_1)	8(P_0), 5(P_1), 4(P_1)	1,2,5,8,7	4	8
8	4 (P_0)	-	1,2,5,8,7,4	4	7
9	4 (P_1)	5(P_1), 1(P_1), 7(P_1)	1,2,5,8,7,4	4	7
10	4 (P_2)	5(P_1), 1(P_1), 7(P_1), 2(P_2), 8(P_2)	1,2,5,8,7,4	4	7
11	4 (P_3)	5(P_1), 1(P_1), 7(P_1), 2(P_2), 8(P_2)	1,2,5,8,7,4 ^e	7	7
12	7 (P_2)	8(P_0), 5(P_1), 4(P_1)	1,2,5,8,7,4 ^e	7	8
13	7 (P_3)	8(P_0), 5(P_1), 4(P_1), 2(P_3), 9(P_3)	1,2,5,8,7,4 ^e	9	8
14	9 (P_0)	-	1,2,5,8,7,4 ^e ,9	9	7
15	9 (P_1)	6(P_1)	1,2,5,8,7,4 ^e ,9	6	7
16	6 (P_0)	3(P_0)	1,2,5,8,7,4 ^e ,9,6	3	9
17	3 (P_0)	6(P_0)	1,2,5,8,7,4 ^e ,9,6,3	3	6

TABLE I

PROTOCOL OPERATION FOR THE EXAMPLE OF SEC.III-D. LISTED FOR EACH CYCLE IS THE STATE INFORMATION FOR THE LEADING NODE.

As expected, energy consumption increases as the number of nodes in the network increases. It can be observed that 50 minutes may be required to set up a network of 81 nodes, whereas around 4 minutes are needed for 9 nodes. Given the complexity and the time required to *physically* deploy a network of 81 or 9 nodes, these figures may be well within acceptable limits.

Fig.6 shows the average number of collisions between reply packets. As discussed earlier, the chances of reply packet colliding are small when the packets are short (2.4 ms in our example). As the number of nodes increases, so does the number of collisions because the network area is kept constant, i.e. the distance between the nodes becomes shorter, and hence the chances of nodes being closer than the critical collision distance increase.




Fig. 6. Average number of reply packet collisions vs. the number of nodes.

As a comparison benchmark, let us consider a broadcast

procedure in which every node transmits its location at the power that is sufficient to span the entire network. The procedure begins with node 1 and moves on after an adequate waiting time needed for all the nodes to receive its broadcast, $\sqrt{2}D/c$. Assuming that each node requires a packet of duration T_P , the network will be discovered in

$$T = N(T_P + \sqrt{2}D/c) \quad (15)$$

and the corresponding energy consumption will be

$$E = NP(\sqrt{2}D)T_P \quad (16)$$

where $P(\sqrt{2}D)$ is calculated according to the expression (9).

Fig.7 compares the broadcast and distributed procedures, showing for each the transmission energy versus the discovery time obtained for a varying number of nodes. We note that the distributed procedure offers a much better energy utilization at the price of longer discovery time. Note, however, that time may be of much less concern than energy in an underwater acoustic network.

Performance of distributed discovery is shown in Fig.7 for two cases, one that uses power control, and another that does not. In the latter case, all transmissions are made at $P_{max} = P(d_{max})$. Power control helps to further reduce the energy consumption, again at the price of somewhat longer time to complete. The energy savings that result from using power control may appear small on the scale of Fig.7, but they amount to several dB.

B. Fading

As discussed earlier, fading may cause the ECP (end-of-cycle packet) to be lost, which will bring the discovery procedure to a premature ending. As a partial remedy, we proposed multiple transmissions of the ECP. Regardless of the number of repetitions, however, there will always be a possibility that some nodes remain undiscovered. To assess

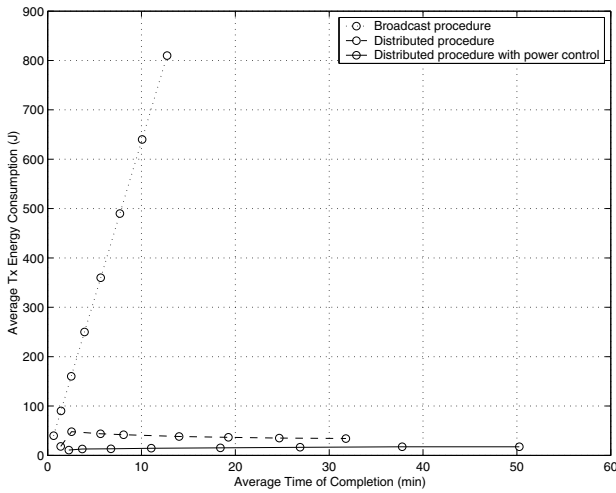


Fig. 7. Transmission energy consumed by the broadcast and distributed procedures. Circles indicate points corresponding to $N=4, 9, 16, 25, 36, 49, 64, 81$ nodes.

the degradation due to fading, we look at the number of nodes discovered as a figure of merit for system *reliability*.

Fig.8 shows the number of nodes discovered vs. the number of nodes actually deployed. Shown in dashed line is the reference case of no fading. Solid and dotted lines correspond to a fading channel with $A/\sigma=30$, with multiple ECP transmissions in two forms, one in which the ECP is transmitted at the same power level as the last discovery packet, and another in which the ECP is transmitted at the next-higher power level (if available).

The first observation to be made is that fading indeed deteriorates the system performance, as evident from the fact that some nodes remain undiscovered. If only a single ECP is transmitted, performance deteriorates beyond an acceptable level (not shown). Multiple ECP transmissions help to recover the performance, notably when they are made at the next-higher power level. Most importantly, we note that a relatively small number of ECP repetitions suffices to recover the performance by a reasonable degree. For example, when 81 nodes were deployed, 60 nodes were discovered using 4 repetitions of the ECP, whereas 75 nodes were discovered using 7 repetitions of the ECP. There also appears to be an effect of diminishing returns when it comes to the number of ECP transmissions, i.e. there is more to be gained by going from 4 to 5 than from 5 to 7, etc. Note that these conclusions are contingent upon the assumption that fading occurs independently between consecutive transmissions, and should be revised for correlated fading.

Fig.9 shows the number of nodes discovered for different values of the ratio A/σ . ECP is repeated either 4 or 7 times, using the next-higher power level. As expected, the performance worsens (fewer nodes are discovered) as the fading strength increases, i.e. as A/σ decreases. Transmitting multiple ECPs remains effective in recovering the performance: with 81 nodes at $A/\sigma=17$, 66 nodes are discovered using 7 ECP repetitions as




Fig. 8. Number of nodes discovered vs. the number of nodes deployed. In case of fading, $A/\sigma=30$.

compared to only 43 nodes using 4 ECP repetitions. Assuming the 7 ECP version, the penalty of fading amounts to 15, 14, 10 and 6 undiscovered nodes out of 81, for A/σ ratio of 17, 20, 24 and 30, respectively. Note that this penalty is conditioned on a given P_{max} , and can further be overcome by investing more power.




Fig. 9. Number of nodes discovered vs. the number of nodes deployed.

Finally, Fig.10 shows the average time it took before the discovery procedure ended (possibly with not all the nodes discovered). This result has to be interpreted in conjunction with Fig.9, i.e. the fact that less time is required under worse fading conditions does not mean that fading helps; on the contrary, it means that fewer nodes were discovered, and, hence, it took less time (and energy) for the procedure to end.

Results similar to those presented here can be generated for varying deployment scenarios (number of nodes, coverage area) and varying fading conditions (best case, worst case expected) to estimate the necessary resources (bandwidth,




Fig. 10. Average time consumption vs. the number of nodes.

energy, maximal power) and protocol adjustments (number of ECP repetitions or similar) needed to achieve a desired performance (reliability, time to complete).

V. CONCLUSIONS

An underwater acoustic network was considered in which the nodes initially have no knowledge of the other nodes' locations, and the network has to be set up autonomously upon deployment. Within this framework, a distributed node discovery protocol was proposed in which the nodes use random access and transmit at minimum power level required to reach a particular neighbor, thus establishing multi-hop connections over a possibly large coverage area where broadcast is not an option within a finite power budget.

Power control is implemented using a small number of distinct levels, and accounts for both the distance-dependent transmission loss and the channel fading. Random access ensures simplicity of implementation with no requirements for global synchronization.

Network initialization proceeds in cycles, each led by a single node. A new leader starts at the lowest power level, increasing it only when necessary, and only up to an available maximum. During discovery, a sign-up sheet is circulated among the leaders. The discovery ends when the sign-up sheet is full, or when all the nodes have exhausted their capabilities (reached maximum power without further progress). Upon completion, each node has built a list of contacts, which includes at least the nearest neighbor, and regular network operation can begin. The key features of the protocol are the simplicity of implementation and the efficient use of power, which are achieved on account of time it takes to complete the discovery.

A simulation analysis was used to quantify the energy consumption and the time to completion. Comparisons with the benchmark case of broadcast initialization clearly demonstrate the benefits of the proposed scheme, whose energy consumption is much lower, while the time to completion,

although increased, stays well within the limits needed to physically deploy a network of underwater nodes (robots, vehicles, floating sensors).

Large-scale fading, modeled as an additional log-normally distributed component on top of a nominal loss, was shown to degrade the system performance, leading to packet loss that may cause a premature ending of the discovery procedure that leaves some nodes undiscovered. Simple adjustments to the discovery protocol (packet repetitions) were shown to recover the performance within limits imposed by the maximum power constraint. Effective performance recovery beyond protocol adjustments is possible by increasing the power budget.

Future work should concentrate on inclusion of more sophisticated fading models that take into account both large-scale and small-scale fading, as well as temporal and spatial fading correlation. These phenomena are bound to affect not only the discovery phase, but also the regular network operation, thus calling for dedicated protocol adjustments on several layers of the network architecture.

ACKNOWLEDGMENTS

This work was supported in part by the NSF grant 0831728 and the ONR grant N00014-09-0700.

REFERENCES

- [1] D. Pompili, I. Akyildiz, "Overview of networking protocols for underwater wireless communications," *IEEE Communications Magazine*, 2009.
- [2] J. Heidemann, W. Ye, J. Wills, A. Syed, Y. Li, "Research challenges and applications for underwater sensor networking," *Wireless Communications and Networking Conference (WCNC)*, 2006.
- [3] M. Chitre, S. Shahabodeen, M. Stojanovic, "Underwater acoustic communications and networking: Recent advances and future challenges," *Marine Technology Society Journal*, 2008.
- [4] B. Peleato, M. Stojanovic, "Distance aware collision avoidance protocol for ad-hoc underwater acoustic sensor networks," *IEEE Communications Letters*, 2007.
- [5] A. Syed, W. Ye, J. Heidemann, "Comparison and evaluation of the T-Lohi MAC for underwater acoustic sensor networks," *IEEE J. Selected Areas in Communications*, 2008.
- [6] N. Chirdchoo, W. Seng Soh, K. Chua, "RIPT: A receiver-initiated reservation-based protocol for underwater acoustic networks," *IEEE J. Selected Areas in Communications*, 2008.
- [7] M. Zorzi, P. Casari, N. Baldo, A. Harris, "Energy-efficient routing schemes for underwater acoustic networks," *IEEE J. Selected Areas in Communications*, 2008.
- [8] D. Pompili, T. Melodia, I. Akyildiz, "Distributed routing algorithms for underwater acoustic sensor networks," *IEEE Trans. Wireless Communications*, 2010.
- [9] JM. Montana, M. Stojanovic, M. Zorzi, "On joint frequency and power allocation in a cross-layer protocol for underwater acoustic networks," *IEEE J. Oceanic Engineering*, 2010.
- [10] G. Alonso, E. Kranakis, R. Wattenhofer, P. Widmayer, "Probabilistic protocols for node discovery in ad-hoc, single broadcast channel networks," *Symposium on Parallel and Distributed Processing*, 2003.
- [11] A. Cornejo, J. Vique, J. Welch, "Reliable neighbor discovery for mobile ad hoc networks," *International Workshop on Foundations of Mobile Computing*, 2010.
- [12] Mc Glynn MJ, SA. Borbash, "Birthday protocols for low energy deployment and flexible neighbor discovery in ad hoc wireless networks," *ACM international symposium on Mobile ad hoc networking & computing*, 2001.
- [13] D. Liu, "Protecting neighbor discovery against node compromises in sensor networks," *IEEE International Conference on Distributed Computing Systems*, 2009.
- [14] S. Vasudevan, J. Kurose, D. Towsley, "On neighbor discovery in wireless networks with directional antennas," *IEEE Infocom*, 2005.

- [15] J. Proakis, M. Stojanovic, J. Rice, "Design of a communication network for shallow water acoustic modems," *Oceans Community Conference*, 1998.
- [16] A. Othman, A. Adams, C. Tsimenidis, "Node discovery protocol and localization for distributed underwater acoustic networks," *International Conference on Internet and Web Applications*, 2006.
- [17] J. Rice, CW. Ong, "A Discovery Process for Initializing Underwater Acoustic Networks," *International Conference on Sensor Technologies and Applications*.
- [18] J. Llor, M. Stojanovic, "A Simulation Analysis of Large Scale Path Loss in an Underwater Acoustic Network," *IEEE OCEANS*, 2011.
- [19] L. Berkhovskikh, Y. Lysanov, "Fundamentals of ocean acoustics," *Springer-Verlag*, 1991.
- [20] M Stojanovic, "On the relationship between capacity and distance in an underwater acoustic communication channel," *ACM Mobile Computing and Communications Review (MC2R)*, 2007.
- [21] S. Basagni, C. Petrioli, R. Petroccia, M. Stojanovic, "Choosing the packet size in multi-hop underwater networks," *IEEE Oceans Conference*, 2010.
- [22] LE. Kinsler, AR. Frey, AB. Coppens, JV. Sanders, "Fundamentals of acoustics," *Wiley*, 1991.
- [23] J.M. Montana, "AUVNetSim: A simulator for underwater acoustic networks," *MIT Sea Grant Technical Report.*, 2008.