WIRELESS COMMUNICATIONS AND MOBILE COMPUTING
Wirel. Commun. Mob. Comput. 2013; **13**:277–295
Published online 14 February 2012 in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/wcm.2206

SPECIAL ISSUE PAPER

A node discovery protocol for ad hoc underwater acoustic networks

Ashish Patil and Milica Stojanovic*

Electrical and Computer Engineering Department, Northeastern University, Boston, MA 02115, U.S.A.

ABSTRACT

Motivated by the advances in acoustic modem technology and the growing number of applications that call for ad hoc deployable autonomous underwater systems (floating sensors, crawlers, vehicles), we address the problem of network initialization upon deployment. A neighbor discovery protocol is proposed, whose goal is to establish communication links over a large area, with a finite power budget that mandates multi-hopping to provide full coverage. The protocol uses random access to eliminate the need for scheduling (i.e., enable system operation without a global clock reference) and power control to ensure that full connectivity is provided using shortest links (i.e., to conserve batteries and prolong the system's lifetime). Transmit power allocation takes into account the acoustic propagation loss, while additional large-scale variation in the average received power is modeled via log-normal fading which is confirmed by experimental observations. System performance is assessed through simulation, by measuring the energy consumption, time to completion, and reliability in the presence of fading. Fading is shown to have a degrading effect on the system reliability, and protocol adjustments are proposed to recover the performance under the constraint on maximum power. The key features of the protocol are simplicity of implementation and efficient use of power. Copyright © 2012 John Wiley & Sons, Ltd.

KEYWORDS

underwater acoustic networks; node discovery; medium access control (MAC); energy efficiency; power control; fading

***Correspondence**

Milica Stojanovic, Electrical and Computer Engineering Department, Northeastern University, Boston, MA 02115, U.S.A.
E-mail: millitsa@ece.neu.edu

1. INTRODUCTION

Research community has taken a considerable interest in the field of underwater acoustic networks, whose growth is motivated by a wide variety of applications that range from environmental sensing by fixed networks of bottom-mounted instruments to exploratory missions by fleets of cooperating autonomous underwater vehicles (AUVs). Ad hoc deployable systems with both mobile and slowly drifting nodes are envisioned for applications such as data gathering and instrument maintenance (e.g., in deep sea oil fields), search and survey missions for both military and commercial applications, and exploratory missions that serve basic sciences. The key technologies that will make such applications possible — vehicles, sensors, and communications — are today mature enough to warrant integration into fully operational systems with a high degree of autonomy.

The focus of our present work is on networking aspects of an ad hoc deployable autonomous underwater system in which acoustic communication is used to form multi-hop links. Acoustic communication is the preferred choice for wireless transmission underwater over distances in excess of several tens of meters. Distances below several tens of meters can be closed by optical links (several meters by radio frequencies in the 10 kHz range), which offer a much greater bandwidth than acoustic links [1]. Optical and radio communications thus offer a *complementary* technique to acoustic communications and are most often considered for fast data download between a sensor and a data mule. In contrast, acoustical links offer a much lower data rate, but they can be implemented over long distances (see e.g., [2] that describes an acoustic link implemented between an 11-km deep ocean trench and a surface ship).

Acoustic communications have seen a rapid development in the past two decades, which resulted in several types of modems that are available both commercially and as research tools. Examples include the Teledyne–Benthos modem [3] and the Woods Hole Oceanographic Institution micro-modem [4]. Because of the system constraints (transducer bandwidth) and the nature of sound propagation which favors low frequencies, the bit rates achieved by acoustic communication are very limited — typically up to several kilobits per second (kbps).

In addition to the very limited bandwidth, acoustic systems are characterized by attenuation that depends not only on the distance but also on the frequency of the signal; a poor quality physical link dominated by the time-varying multipath propagation and severe Doppler effects and long delays caused by the low speed of sound (nominally 1500 m/s). These characteristics distinguish underwater acoustic systems from the better known terrestrial radio systems and have a profound implication not only on the design of physical layer techniques [5] but also on the network protocols [6]. Respecting the engineering aspects of system design and the physics of acoustic propagation are of paramount importance in setting the network optimization criteria. Perhaps, the most important fact to be kept in mind is that underwater networks are required to operate in extremely harsh environments, which mandate expensive devices that need to withstand the pressure and the forces of waves and currents and whose batteries are difficult to recharge once deployed.

The fundamental aspects of underwater acoustic networking have been highlighted in several recent publications, for example, [6–8]. In the past years, there have been major developments on the medium access control (MAC) layer, with focus on protocol 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. Examples of such protocols include [9,10] and [11]. Simultaneously, work has been active on the higher layer protocols, in particular routing. Several protocols have been proposed that explicitly take into account acoustic system issues, for example, [12–14].

Cross-layer design plays an important role in acoustic systems, both between adjacent network layers and between adjacent system functions. Examples of cross-layer design include power control integration with MAC and routing [14]; packet size selection for optimizing throughput/energy in conventional settings [15] as well as in network-coding settings [16] with half-duplex acoustic modems and topology control for energy-efficiency in fixed networks [17]. An example of cross-*function* design is integration of localization/navigation with acoustic communications. The use of acoustic modems as traveling beacons that help to localize AUVs by measuring the relative propagation delays in a network is described in Reference [18]. Such localization may in turn support geographical routing protocols (e.g., [14]) in which location information is used to find the best routes through the network.

Although there are no routinely operational underwater networks at this time, experimental networks are gaining momentum [19], and research is producing a solid base of protocols that are suitable for poor-quality, interference-prone, delay-challenged acoustic channels, and are ready for testing in the field. In this paper, we draw on the availability of such tools in service of the powerful vision of ad hoc deployable underwater networks and study the problem of network initialization.

2. THE NODE DISCOVERY PROBLEM

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 the their neighbors' IDs or locations upon deployment — an assumption on which a typical MAC and routing protocol suites rely. The network's first task is thus to establish the communication links. The goal of the discovery procedure is to inform all the nodes of their neighbors' identity and the associated power level needed to establish full network connectivity. Only after each node has discovered its neighbors and established the paths to them can a regular network operation begin. The information gathered about the network topology during initialization can be used to build routing tables, which can later be dynamically updated, and it can also be passed on to a gateway or a command center if such exists.

Node discovery has been extensively investigated for radio networks, but work in the field of underwater acoustic networks is extremely scarce. A specific neighbor discovery procedure for an underwater acoustic network was first discussed in [20]. This reference proposed a protocol based on polling by a master node in a centralized configuration. The channel access was regulated through code division-multiple access. Reference [21] proposed a similar neighbor discovery procedure. In Reference [19], an experimental deployment is described in which initial neighbor discovery is performed via a master node.

Some of the works on node discovery in terrestrial networks include References [22,26]. Reference [22] describes a probabilistic protocol for node discovery. In Reference [23], a reliable neighbor discovery layer is defined for mobile ad hoc networks. Birthday protocols [24] are a family of probabilistic protocols for static ad hoc networks, which attempt to conserve energy during the deployment phase while increasing the probability of neighbor discovery during the discovery phase. Reference [25] presents a secure neighbor discovery protocol for a compromised node in a static network. In Reference [26], neighbor discovery in static networks with directional antennas is presented.

In this paper, we propose a decentralized node discovery procedure for an underwater acoustic network. The nodes are assumed to be deployed over a possibly large area that cannot be spanned in a single hop within the constraints of a finite power budget. The nodes operate in a distributed manner, that is, without a central station or a-priori regulation of the channel access. We consider a random access environment, that is, a network in which there is no a-priori division of the available resources such as time, frequency, or code-division multiple access (TDMA, FDMA, CDMA) and no scheduling. Random access is advantageous from the viewpoint of simplicity of implementation which requires no global synchronization for scheduling.

 Wirel. Commun. Mob. Comput. 2013; **13**:277–295 © 2013 John Wiley & Sons, Ltd.
DOI: 10.1002/wcm

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, that is, it must guarantee, within the power constraint, a path through the network between any two participating nodes.† The following assumptions are made:

- There is a finite number of nodes, N.
- The nodes know the number N.
- Nodes know its location with a certain accuracy.
- Nodes can vary their transmission power 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, just that the local neighborhood topology does not change faster than the nodes can learn it. Given a typical AUV speed of a few meters per second and a temporary confinement area of several hundred square meters, this is a reasonable assumption.

Power control is implemented in discrete levels and accounts for both the distance-dependent transmission loss and the channel fading. Power control is implemented in an open loop fashion so that a node decides itself which power level to use for transmission rather than being instructed by another node.

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 [15], 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. 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 commence.

Figure 1 illustrates a newly deployed network which lacks structure, whereas Figure 2 shows the network after completion of discovery. All the nodes in Figure 2 know their neighbors and the power level required to reach them.

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 in a similar manner as it initiated the original discovery.

The rest of this paper is organized as follows. Section 3 is devoted to describing the fundamental mechanisms of acoustic propagation. We describe both the basic propagation loss and the large-scale multipath fading effects. In Section 4, we present the node discovery protocol. This

†A node that cannot be reached by another node at the maximal power level is considered not to be participating.

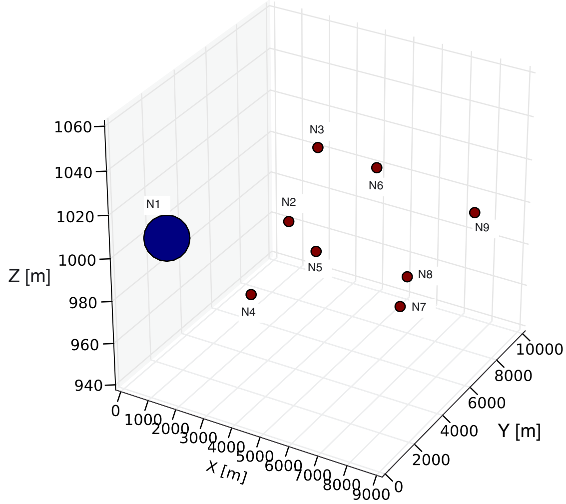

Figure 1. A newly deployed network has no structure, that is, no connections between the nodes.

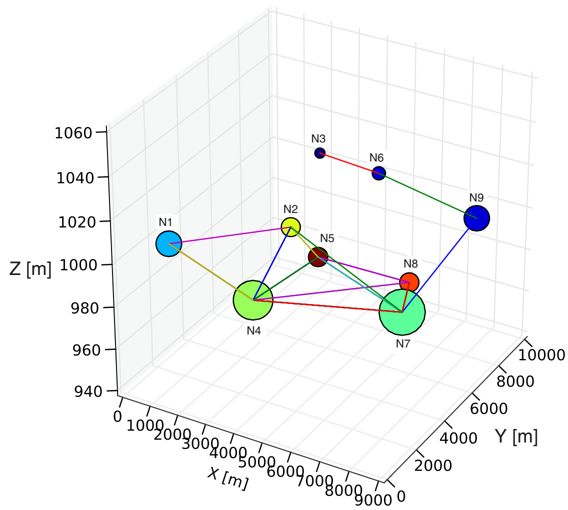

Figure 2. Node connections are established during discovery.

section contains an outline of the network topology, signaling format, and channel access as well as a detailed illustration of the protocol operation. Performance analysis is presented in Section 5, where simulation results are given for both the ideal case with no fading and for the realistic case of fading, and comparisons are made with the benchmark case of broadcast initialization. Finally, conclusions are summarized in Section 6.

3. ACOUSTIC PROPAGATION

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. Several modeling stages can be used to assess the received signal properties in varying detail. The first stage addresses the basic, deterministic propagation loss that occurs because of energy spreading and absorption, whereas the later stages address finer aspects of multipath propagation (reflection, refraction, scattering, etc.) that contribute to the signal strength variation around the nominal value predicted by the basic loss. Signal strength varies as transmitter, and receiver assume particular locations in a given channel geometry, or as the channel geometry and the propagation conditions vary in time. Some of these variations can be modeled in a deterministic manner, whereas other appear to be random.

Generally speaking, there are two types of channel models: small-scale models that target the *instantaneous* signal power and are of interest for receiver-end signal processing (physical layer) and large-scale models that target the *average* received signal strength (average over some local interval of time or area of displacement) and are of interest for power allocation at the transmitter (top-level system design). The latter type of modeling is of interest to our present study. In the following section,, we first discuss the basic propagation loss then move on to large-scale modeling.

3.1. Basic loss

The attenuation, or path loss, experienced by an acoustic signal of frequency f as it travels over a distance d in an unobstructed medium is given by [27],

$$A(d, f) = A_0 d^k \mathrm{a}(f)^d \qquad (1)$$

where A_0 is a scaling constant, k is the spreading factor, and $\mathrm{a}(f)$ is the absorption coefficient. The basic loss has two components: spreading loss, described by the factor d^k and the absorption loss, described by the factor $\mathrm{a}(f)^d$. The spreading factor k models the geometry of spreading, which can range from cylindrical (k=1) to spherical (k=2, as in radio systems). The absorption loss occurs because of acoustical energy conversion into heat. The absorption coefficient $\mathrm{a}(f)$ increases rapidly with frequency, as shown in Figure 3. Absorption thus confines the acoustic frequencies to low values. For example, an acoustic system designed for transmission over 1 km could operate in 10 kHz of bandwidth centered at 10 kHz, that is, between 5 and 15 kHz. It is interesting to note that although the acoustic bandwidth is very small compared with that of radio systems, an acoustic communication system is in fact wideband in the true sense of the word — its bandwidth is not negligible with respect to the center frequency. This fact has important implications on the design of physical layer techniques, notably synchronization and array processing.

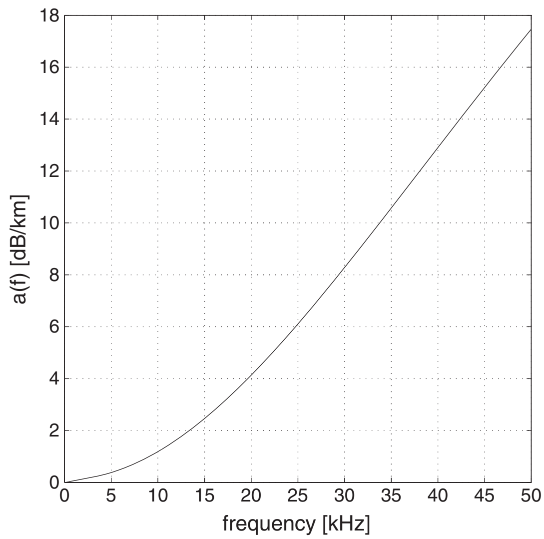

Figure 3. The absorption coefficient $\mathrm{a}(f)$.

Background noise in an underwater acoustic system comes from turbulence, shipping and breaking waves, and also includes thermal noise. Although the background noise is often described as Gaussian, it is not white. The power spectral density of the background noise can be approximated as [28]

$$N(f) \approx N_0 f{-}\eta \qquad (2)$$

where N_0 is a constant that can be measured for a particular system location, and the factor η models the power spectral density decay, which usually occurs at 18 dB/decade. The noise level decreases with frequency, thus creating an opposite effect from absorption. As a result, there exists an optimal center frequency to be used for a given distance. The optimal center frequency decreases with distance and so does the bandwidth that can be allocated around it. An interested reader is referred to [28] for details.

From the networking point of view, the fact that the available bandwidth decreases with distance implies that a higher bit rate will be available if transmission is organized over multiple short hops instead of one long hop. Multihopping thus appears not only as a power-saving strategy (which is the case in any wireless network) but also as a bandwidth-increasing strategy. The problem, of course, is

Wirel. Commun. Mob. Comput. 2013; **13**:277–295 © 2013 John Wiley & Sons, Ltd.
DOI: 10.1002/wcm

that by increasing the number of hops, the level of interference may increase, and packet collisions may become more likely. However, this does not have to be so. Namely, as shorter hops support higher bit rates, data packets containing a given *number* of bits will have shorter *duration*, and the chances of collision will be reduced. This fact speaks further in favor of multi-hopping in an underwater acoustic scenario.

The signal power, received at a distance d in response to a transmit power P_T allocated uniformly across a signal bandwidth B centered at frequency f_c is given by

$$P_R = \frac{P_T}{B}\int_{f_c-B/2}^{f_c+B/2} A^{-1}(d, f)df \equiv P_T \cdot \overline{G}(d) \quad (3)$$

where $\overline{G}(d)$ represents the nominal channel gain (inverse of attenuation). The noise power is given by

$$P_N = \int_{f_c-B/2}^{f_c+B/2} N(f)df \quad (4)$$

and the signal-to-noise ratio (SNR) is obtained as

$$\mathrm{SNR} = \frac{P_R}{P_N} \quad (5)$$

The SNR is used as a principal figure of merit for determining the transmit power necessary to achieve a certain level of performance. In particular, if it is required that the SNR be above a pre-specified threshold SNR_0, the transmit power needed to close a distance d is determined as

$$P_T \geq P(d) \quad (6)$$

where

$$P(d) = \mathrm{SNR}_0 \frac{P_N}{\overline{G}(d)} \quad (7)$$

We will use the aforementioned definition of the necessary transmit power $P(d)$ later when we discuss power control.

3.2. Multipath

Although the basic propagation loss describes energy spreading and absorption, it does not take into account the specific system geometry and the resulting multipath propagation. Multipath propagation occurs because of various phenomena including surface-bottom reflections and ray bending, which is mostly notable in areas where the speed of sound changes with depth. The first approach to modeling multipath propagation is a deterministic one, which provides an exact solution for the acoustic field strength in a given system geometry with a given sound speed profile. Because of computational complexity, approximate solutions are often used instead of the exact. An approximation that is suitable for frequencies of interest to acoustic communication systems is based on ray tracing (a popular package is available online [29]). Figure 4 illustrates a ray-trace obtained for a transmitter placed inside the circle to the left. Lighter colors in this figure indicate locations of higher received signal strength. If one places a receiver at some distance away from the transmitter in this field, one notices that the signal strength will vary depending upon the *exact* location. In other words, two receivers (e.g., two circles to the right) placed at the *same* distance away from the transmitter, may experience propagation conditions that are quite different. In this example, shadowing is caused by ray bending in deep water. In shallow water with constant sound speed, signal strength can be calculated using simple geometrical considerations. Alternating constructive/destructive combining of multiple reflections will now form pockets of strong/weak signal reception. In either case, multipath effects make the signal strength location-dependent, i.e., different from the value predicted by the basic propagation loss (1). The signal strength can still be calculated using an expression similar to (3) but with appropriate adjustments made to include the full channel transfer function $H(d, f)$ in place of the ideal, unobstructed-single-path's $A^{-1}(d, f)$.

3.2.1. Random effects.

If only a deterministic model is employed to assess the multipath effects, the channel transfer function $H(d, f)$ can be completely specified. Experimental observations, however, provide ample proof that the situation is not so simple. It is well-known that propagation conditions vary both with time and with small displacements of the transmitter/receiver. In other words, a transmitter/receiver pair separated by the same distance d but in two different locations may experience different channels. Displacements may occur because of either intentional motion (AUVs) or unintentional motion, as free-floating devices move with currents and waves. The surface also moves with tides and

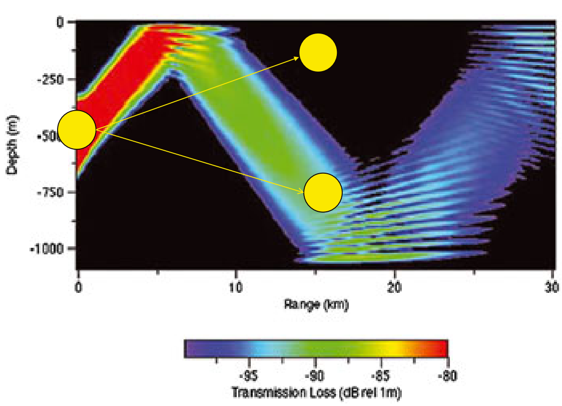

Figure 4. A ray trace shows areas of strong/weak signal reception. Sound reflection and refraction result in multipath propagation, which favors some locations while placing others in a shadow.

waves, creating a time-varying propagation environment. The effects of changing propagation conditions are evident as changes in the received signal strength that are too tedious to assess analytically and are instead modeled as random processes.

A large-scale propagation model that we will use in this work gives the received signal strength as

$$P_R = P_T \cdot G(d) \tag{8}$$

where the gain $G(d)$ is now treated as a random variable.

Figure 5 shows an ensemble of experimentally recorded gains expressed on the logarithmic scale (in dB),

$$g(d) = 10 \log_{10} G(d) \tag{9}$$

We note that the gain exhibits a decreasing trend with distance as expected because of energy spreading and absorption, but we also note variations around this decreasing trend. This observation motivates us to model the gain as

$$g(d) = \overline{g}(d) + x \tag{10}$$

where $\overline{g}(d)$ is a mean value, and x is a random variable.

Figure 6 shows the histogram of the residual x. The values of x are obtained by removing an estimated mean (solid curve in Figure 5) from the measured gain. The resulting histogram is shown along with a Gaussian probability density function of zero mean and variance σ^2, which is calculated from the measurements.

Encouraged by the similarity between the measured and hypothesized probability density functions of Figure 6, as well as by recent findings [30], [31] which testify to a similar effect, we propose a log-normal model for the large-scale gain variation. In this model, the gain (9) is specified by the mean $\overline{g}(d)$, which can be predicted using a deterministic model for a given *nominal* channel geometry and a random component x, which is distributed according to $\mathcal{N}(0, \sigma^2)$, where σ^2 is independent of the distance (at least for some range) and can be estimated from measurements.

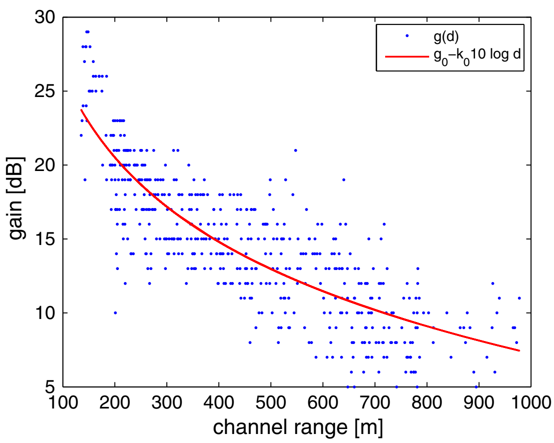

Figure 5. Gain (normalized) versus transmission distance. Dots show measured values; solid curve shows an estimated trend (a first-order logarithmic-scale polynomial fit to the ensemble mean at each distance yields k_0=1.9).

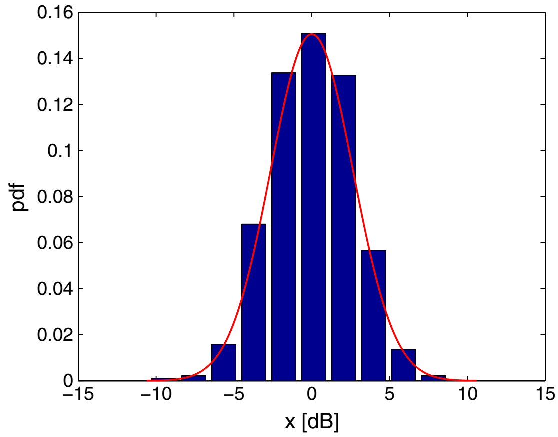

Figure 6. Histogram of the measured deviation x and the theoretical probability density function of a zero-mean Gaussian random variable with σ^2=6.7 dB.

Specifically, we will use this model to assess the protocol performance in conditions of fading. For the sake of generality, we will not focus on a particular multipath geometry but will instead use the basic path loss (1) to model the mean value $\overline{g}(d)$. We will thus focus on the effects of large-scale fading and investigate their impact on the system performance using different values of the variance σ^2.

4. THE DISCOVERY PROTOCOL

The discovery begins arbitrarily with node 1, who leads the first cycle. A new leader starts at the lowest power level, increasing it only when necessary and only up to an available maximum. As the discovery proceeds, a sign-up sheet is filled, eventually indicating that all the nodes that are within each other's reach have been accounted for. Upon completion, each node will also have built a list of contacts, which includes at least the nearest neighbor. The key features of the protocol are simplicity of implementation and efficient use of power.

In this section, we begin by describing the network topology and the power control mechanism. We then specify the signaling format and outline the steps that define a discovery cycle. We comment on the channel access and discuss the protocol adjustments needed to overcome the deleterious effects of fading. We finally illustrate the protocol operation using an example and provide a complete summary of the protocol functions.

Wirel. Commun. Mob. Comput. 2013; **13**:277–295 © 2013 John Wiley & Sons, Ltd.
DOI: 10.1002/wcm

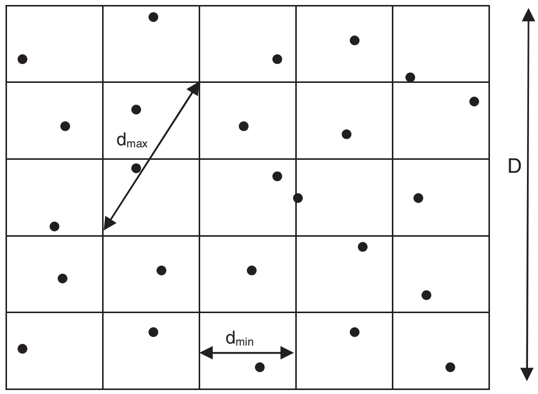

Figure 7. Nodes randomly located within a grid.

4.1. 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 Figure 7. Such a scenario is a representative of a mission in which a group of nodes (crawlers, vehicles) are dispatched to map a given region, and each is assigned a smaller individual area. Accordingly, we assume that a node is placed randomly within its cell.

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

$$d_l = d_0 + l\Delta d, l = 0, \ldots, L-1 \quad (11)$$

Given a total of N nodes, the minimum and the maximum distance (see Figure 7) are set to

$$d_0 = d_{min} = \frac{D}{\sqrt{N}} = \frac{1}{\sqrt{\rho}} \quad (12)$$

and

$$d_{L-1} = d_{max} = d_{min}\sqrt{5} \quad (13)$$

where ρ denotes the node density.[‡] Given a desired number of power levels, the step Δd is determined as

$$\Delta d = \frac{d_{max} - d_{min}}{L-1} \quad (14)$$

We note that if one were to choose a uniform step in power, say several dB, such a choice would differ little from the one based on a uniform step in distance [14].

Transmission power required to provide a target SNR_0 at a particular distance d_l is now obtained from the expression (7). We denote this power by $P_l = P(d_l)$.[§]

[‡]We assume without loss of generality that N is a square number, that is, N=4, 9, 16, 25, 36, and so on.

[§]Conversion from acoustical power expressed in in dB re μPa^2 to electrical power expressed in dB re W is performed by subtracting 172 dB [32].

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.

4.2. Signaling format and the discovery cycle

The signaling format is specified by the packet structure and the power level. The discovery packet contains the following information: [type: D, source, destination, power level]. Because 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 level and repeats the discovery packet. The amount of time it takes to complete a transaction at power level l, that is, the amount of time a leader will wait before switching to the next state is

$$T_l = (T_P + \tau_l) + (T_R + \tau_l) + T_G \quad (15)$$

where T_D is the duration of discovery packet, T_R is the duration of 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 (ECP). The sign-up sheet is enclosed in this packet. It consists of N fields, each corresponding to one of the nodes. A leader signs up by putting its ID in an appropriate field.

The outgoing leader also selects the next leader from its list of contacts and includes this information into the ECP. 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 choice is that shorter links require less power (although they may prolong the discovery).

The total amount of time during which one node holds a leadership, starting at power level 0 and passing it on at some level l, is

$$T_l = \sum_{i=0}^{l} t_i + T_E + \tau_l \tag{16}$$

This expression can be used to roughly estimate the bounds on the total time (and energy) needed to initialize a given network under ideal propagation conditions. A lower bound will correspond to the best case in which discovery proceeds directly from node to node, and each node uses only the lowest power level. An upper bound will correspond to the worst case in which each node cycles through all the power levels and cycles may need to be repeated. These bounds will represent only rough estimates because the actual time and energy will include retransmissions needed to make up for the packets lost in both collisions and fading, which we discuss next.

4.3. 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. (Because only one node at a time can be the leader, there are no collisions among the discovery packets or the ECPs.) However, 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 about 2 m. Hence, the chances that reply packets will collide at the leader are small. Nonetheless, the protocol has to account for such a possibility.

4.3.1. 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 initiates a collision recovery procedure. Under this procedure, 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 because 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 retransmitted. This time, the nodes will not transmit immediately. Instead, each node will wait for a random back-off time and retransmit only then. As in any back-off procedure, the additional random delay will help to reduce future collisions.

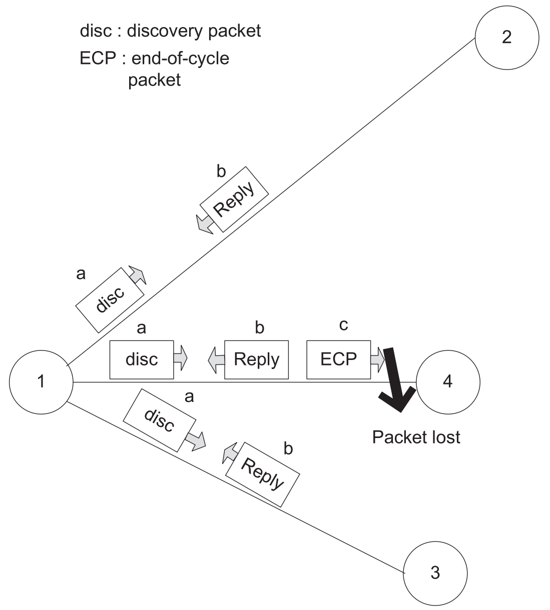

Figure 8. End-of-cycle packet can be lost due to fading, causing the discovery procedure to stop. A possible remedy for this situation is to transmit the end-of-cycle packet multiple times (if one transmission experiences fading, chances are that another may not).

4.3.2. End-of-cycle.

For the discovery protocol to proceed correctly, it is necessary that the 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. Figure 8 illustrates such a situation.

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. Although 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 ECP is received. In particular, the protocol can mandate that the ECP be acknowledged, or that it be repeated a sufficient number of times. Because 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 at the same power 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). In our discussion of results in Section 5, we will assume independent fading realizations from one transmission to the next and show that a small number of immediate repetitions then suffices to recover the system performance to a reasonable degree. Simply stated, if one ECP is lost, the chances of it being lost *again* are small if each new retransmission sees an independent channel realization.

Although many different protocol adjustments are possible to guard against fading, the simple end-of-cycle repetition suffices to illustrate the principles. We also note that the system can take advantage of *implicit* acknowledgments, which are inherent to the broadcast medium. In other words, upon overhearing the *next* cycle, the leader of the *former* cycle will know that its ECP has been received. Conversely, if it has not heard anything in a pre-specified amount of time, the former leader can re-initiate the discovery procedure. Such protocol details are a matter of particular system design and of no concern at the moment.

4.4. Performance illustration

Referring to Figures 1 and 2, in this section, we illustrate the steps of the discovery procedure. The steps are organized according to the leader who is in charge of the indicated cycles. 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 others' reach are accounted for by the end of the discovery phase. In doing so, the goal is to limit the total power consumption, that is, not all the connections have to be discovered, but full connectivity has to be provided. Table I provides a reference summary of the system state at the end of each cycle.

Cycles 1, 2: Node 1 forms a discovery packet and transmits at the lowest power level P_0. In this example, it does not find any nodes. Hence, it retains the leadership for the next cycle and transmits at the next power level. Let us say that at power level P_1 (cycle 2), nodes 2 and 4 are reached. The leader, in addition to storing the replying nodes' IDs, also remembers the power level at which those nodes were reached. The leadership is then passed on to the closest node in the list of contacts, node 2. The network connections after cycle 2 are shown in Figure 9. In the figure, the size of a node is proportional to the energy that it has consumed while transmitting.

Cycle 3: After becoming a leader, node 2 transmits the discovery packet and finds node 5 at power level P_0. Node 2 then chooses node 5 to be the leader for cycle 4, because it is the only node in the list of contacts and has not been a leader before. In addition to remembering which nodes were discovered, each node will also remember who

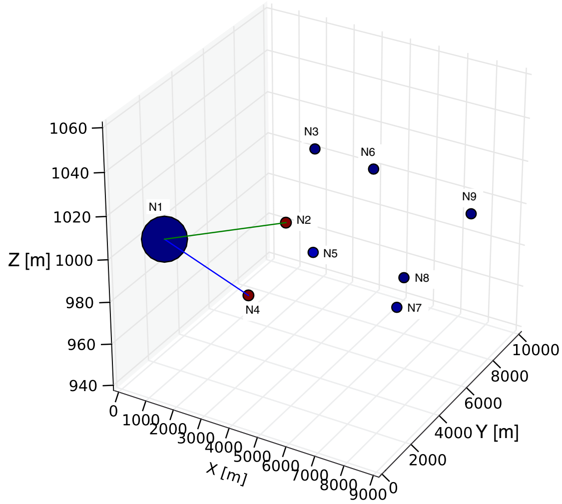

Figure 9. Network at the end of cycle 2 led by node 1 $(N1)$. List of contacts: $N1 = [N2(P_1), N4(P_1)]$.

Table I. Protocol operation for the example of Section 4.4. Listed for each cycle is the state information for the leading node.

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

selected it to be the leader. Thus, node 2 will remember that it was selected by node 1. Note that it is the original elector that a node remembers, not the last node that returned the leadership (we will see later how the leadership can be returned).

Cycle 4: On assuming leadership, node 5 transmits the discovery packet and finds nodes 2 and 8 at power level P_0. Node 2 has already been a leader, so node 5 passes the leadership on to node 8.

Cycle 5: Node 8 is now the leader and it transmits the discovery packet to find nodes 7 and 5 at P_0. The leadership is then forwarded to node 7, because it is the closest node in the list of contacts and has not been a leader before.

Cycles 6, 7: Node 7 finds node 8 at power level P_0 (cycle 6), but node 8 has already been a leader. Consequently, node 7 proceeds to increase the power looking for a node which has not already been a leader. In cycle 7, it finds nodes 5 and 4 at P_1. Node 7 passes the leadership on to the closest node which has not been a leader, which is node 4.

Cycles 8, 9, 10, 11: On becoming the leader, node 4 finds nodes 1, 5, and 7 using power level P_1, but they have already been leaders. Thus, node 4 increases the power to the next level. It finds nodes 2 and 8 using P_2, but they too have been leaders before. Node 4 thus proceeds to test the next power level, P_3. Let us assume that this is the maximal power level but that still no new node is discovered. Because it has reached the maximal power level, there is nothing more that node 4 can do. However, the sign-up sheet is not full, and the discovery procedure should go on. To continue the discovery, the leader will give the leadership back to its return address (a node from whom it received the leadership). Thus, node 4 gives the leadership back to node 7. At the same time, node 4 adds an end mark next to its name on the sign-up sheet. The end mark indicates that a node has exhausted its capabilities. In other words, it indicates a dead end, as there is no point in asking this node to be a leader again because it would yield no results. The scenario after cycle 11 is shown in Figure 10.

Cycles 12, 13: Node 7 has been a leader before. When it becomes a leader again, it first checks its list of contacts to see if it contains any node that has not been a leader before. If it finds any such node, then it will directly transmit the ECP to that node, giving it the leadership. In this particular case, there is no such node on the list. Node 7 has previously transmitted at P_1; hence, by transmitting at a higher level, it may find a new node. It thus transmits the discovery packet at power P_2 but finds no new nodes. Finally, at power P_3, it finds nodes 2 and 9. Node 2 has been a leader before; hence, it passes the leadership to node 9.

Cycles 14, 15: Here node 9 is the leader. It finds node 6 at power P_1 in cycle 15. Node 9 passes the leadership on to node 6.

Cycle 16: Node 6 finds node 3 at power P_0 and passes the leadership on to it. This will be the last node in our example, but the discovery procedure does not end until node 3 has put its ID on the sign-up sheet.

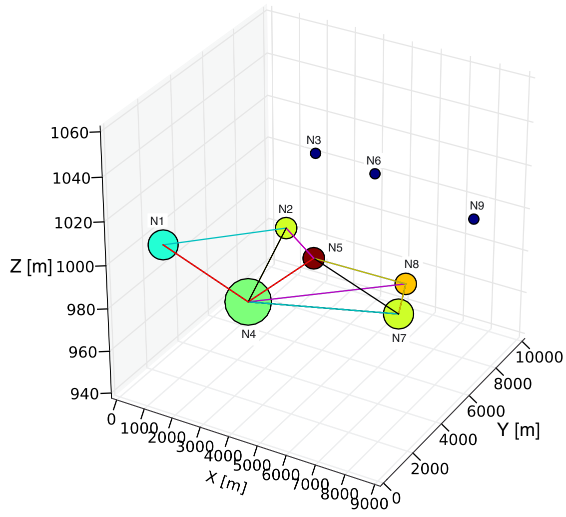

Figure 10. Network at the end of cycle 11 led by node 4 ($N4$). List of contacts: $N4 = [N5(P_1), N1(P_1), N7(P_1), N2(P_2), N8(P_2)]$.

Cycle 17: Although the sign-up-sheet is full, node 3 starts transmitting at the lowest power level according to the protocol. By starting at the lowest power level, node 3 is guaranteed to find all its nearest neighbors, if any exist. Node 3 thus finds node 6 at P_0. It puts its ID on the sign-up sheet, which is now full and the procedure ends.

In general, 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, that is, within the given power constraint.

Once the discovery procedure has ended, all the nodes in the network can begin regular operation. At this time, that is, at the end of cycle 17 in our example, 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).

4.5. Protocol summary

The complete protocol is formally defined by the block diagrams of Figures 11 and 12. The block diagram of Figure 11 summarizes the steps that should be followed

Wirel. Commun. Mob. Comput. 2013; **13**:277–295 © 2013 John Wiley & Sons, Ltd.
DOI: 10.1002/wcm

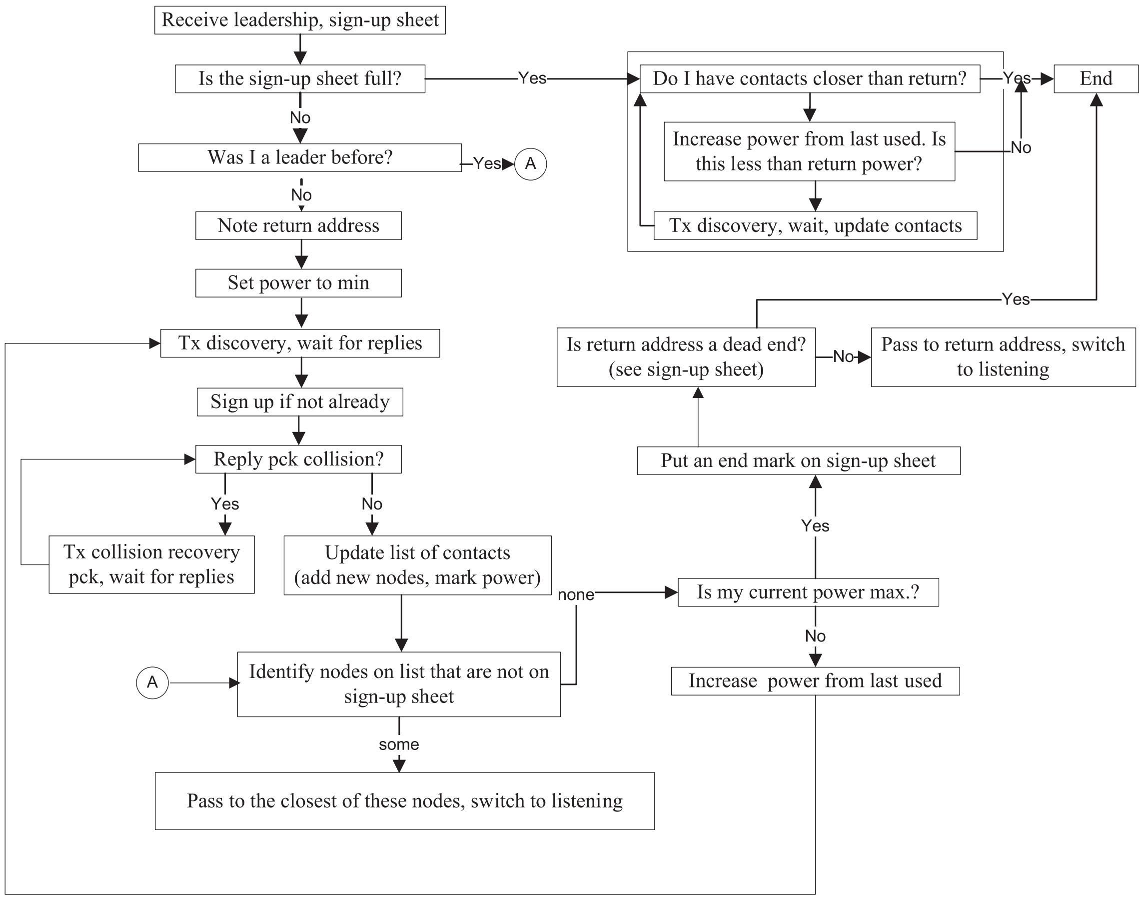

Figure 11. Functions performed by a node acting as leader.

by a node when it is acting as a leader. The block diagram of Figure 12 summarizes the functions followed by a node that is not a leader (this node is either listening passively, waiting to detect a discovery packet, or it is actively engaged in replying to the leader).

5. PERFORMANCE ANALYSIS

To study the performance of the protocol, we conducted a simulation analysis using AUVNetSim, a discrete event simulator written in Python [33]. In this section, we study two cases: an ideal case in which there is no fading and a case with fading. The protocol performance is assessed by measuring the energy consumption and the time to completion. Comparisons are made with the benchmark case of broadcast initialization to quantify the benefits of distributed discovery.

5.1. Simulation parameters

The following scenarios were considered:

- a square area of side D = 10 km
- center frequency f_c = 10 kHz
- bandwidth B = 10 kHz
- target SNR_0 = 25 dB
- background noise level N_0 = 50 dB re $\mu\mathrm{Pa}^2/\mathrm{Hz}$[¶]
- number of power levels L = 4
- number of nodes N=4, 9,16,25,36,49,64,81
- $d_{min} = D/\sqrt{N}$=5, 3.34,2.5,2,1.67,1.25,1.11 km
- fading parameter A/σ=17,20,24,30
- reply packet size: 24 bits
- number of end-of-cycle packet transmissions: 1,4,5,7

The results are averaged over 1000 random deployments of nodes (a node is placed uniformly with its cell, independently of other nodes).

5.2. Ideal case (no fading)

Figure 13 shows the average total energy consumption versus the discovery time obtained for a varying number

[¶]Note that this level depends on the particular system location and can vary by several tens of dBs. However, considering a higher noise level will not alter the general performance *trends*, only the absolute values.

DOI: 10.1002/wcm

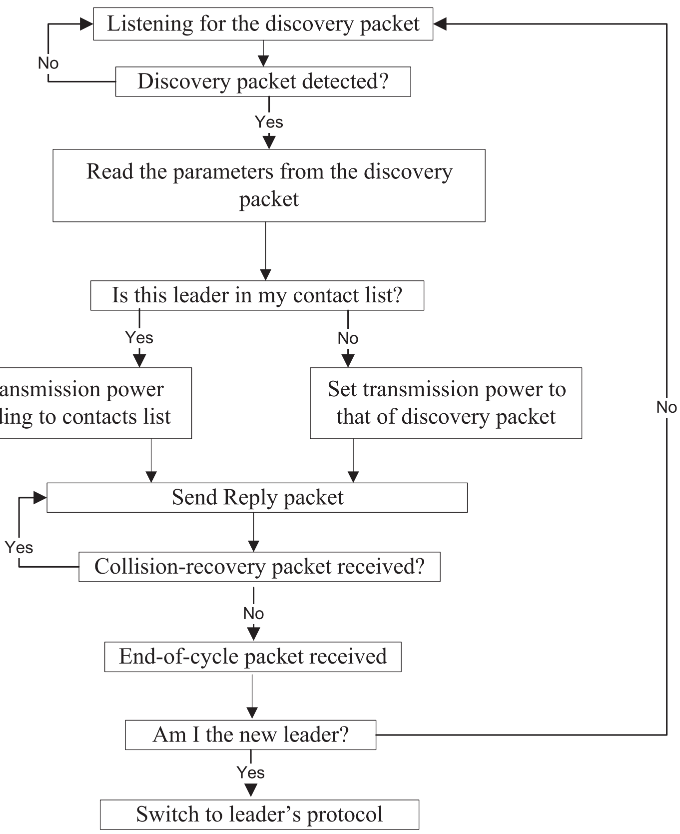

Figure 12. Functions performed by a node not acting as a leader.

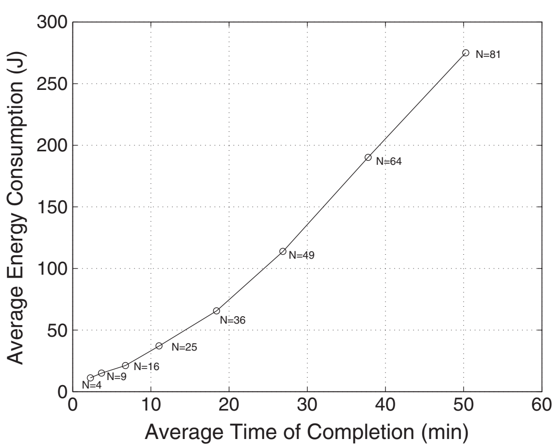

Figure 13. Total energy (average) versus the time it takes to complete the discovery.

of nodes. The total energy is the sum of energy spent on transmission, reception, and idle listening. As expected, energy consumption increases as the number of nodes in the network increases. It can be observed that 50 min may be required to set up a network of 81 nodes, whereas around 4 min are needed for nine nodes. Given the complexity and the time required to physically deploy a network of 81 or nine nodes, these figures may be well within acceptable limits.

Figure 14 shows the average energy consumed for reception versus the number of nodes. Unlike the transmission energy, the reception energy does not depend on the distance between the nodes but only the number of nodes in the network. The reference reception power is set to 0.1 W.

Figure 15 shows the average number of collisions that are taking place between reply packets. As discussed earlier, the chances of reply packet colliding are small because the packets are short. Assuming a bit rate equal to the bandwidth, the reply packet lasts for T_R=24/B=2.4 ms, and two such packets, transmitted immediately in response to the discovery packet, will collide if the underlying distances to the leader differ by less than $T_R \cdot c/2$=1.8 m.

5.3. Comparison with broadcast initialization

As a comparison benchmark, we 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

Wirel. Commun. Mob. Comput. 2013; **13**:277–295 © 2013 John Wiley & Sons, Ltd.
DOI: 10.1002/wcm

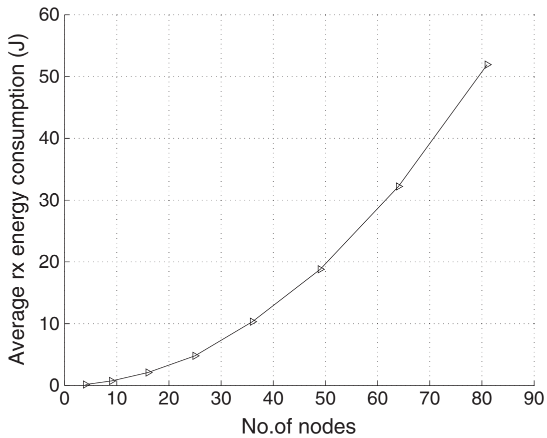

Figure 14. Average receiving energy versus the number of nodes.

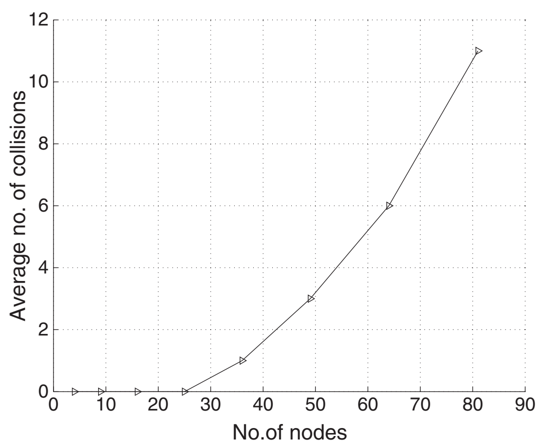

Figure 15. Average number of reply packet collisions versus the number of nodes.

waiting time needed for all the nodes to receive the broadcast, that is, $\sqrt{2}D/c$. Assuming that each node requires a packet of duration T_P, the network will be discovered in

$$T = N(T_P + \sqrt{2}D/c) \tag{17}$$

Because every node's transmission will have reached every other node by the end of the waiting time, there is no need for more than one round in this case. The power needed by each node to broadcast is $P(\sqrt{2}D)$, which can be computed using expression (7). The corresponding energy consumption is

$$E = NP(\sqrt{2}D)T_P \tag{18}$$

Figure 16 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.

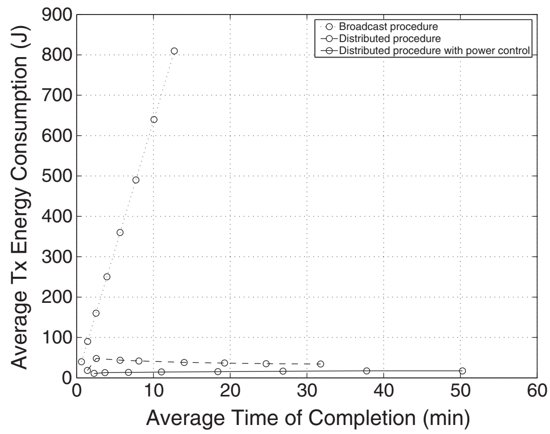

Figure 16. Transmission energy consumed by the broadcast and distributed procedures. Circles indicate points corresponding to N=4, 9, 16, 25, 36,49, 64, 81 nodes.

Note, however, that time may be of much less concern than energy in an underwater acoustic network.

Performance of distributed discovery is shown 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 clearly 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 Figure 16, but they amount to several dB.

5.4. Fading

Existing literature on the subject of underwater acoustic networks almost exclusively analyzes the network performance under ideal conditions. In practice, however, it is well-known that underwater acoustic channels exhibit time-variation (fading) and it is only plausible to expect the network performance to be affected. Specifically, two types of questions arise in the context of fading channels. The first is 'What is the effect of fading on the network performance, i.e. how serious is the degradation'? The second is 'What can be done to remedy the situation'? To answer these questions, we have included the fading model of Section 3.2.1 into our simulation analysis. Specifically, we conducted simulations for varying fading parameter A/σ, which represents the ratio of the absolute mean value of the gain $\overline{g}(d)$ to the standard deviation σ of the random component x. This section discusses the results that we obtained.

5.4.1. Premature ending / undiscovered nodes.

As discussed in Section 4.3.2, fading may cause the ECP to be lost, which will bring the discovery procedure to a premature ending. As a remedy, we proposed multiple transmissions of the ECP. Note, however, that regardless of the number of repetitions, in the presence of fading, there

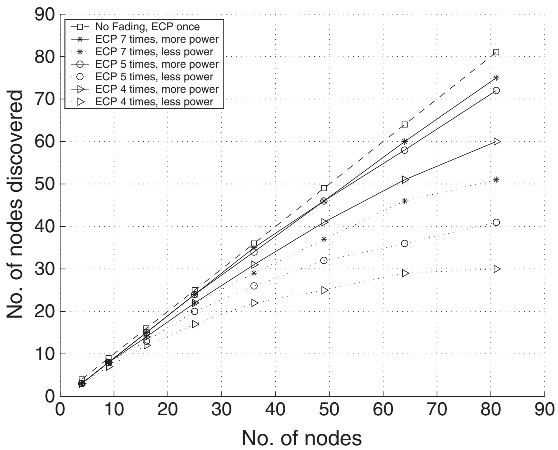

Figure 17. Number of nodes discovered versus the number of nodes deployed. In case of fading, $A/\sigma=30$.

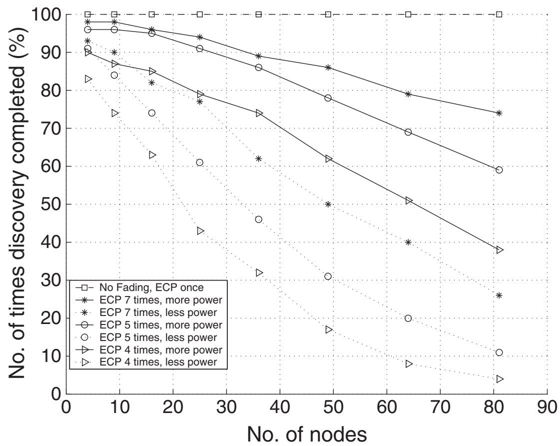

Figure 18. Percentage of times discovery procedure was successfully completed (all nodes discovered) versus the number of nodes in a fixed area.

will always be a possibility that some nodes will remain undiscovered. In what follows, we will look at the number (percentage) of nodes discovered as a figure of merit for system *reliability*. We will also look at two types of ECP repetitions: 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). Simulation results are now averaged both over random node deployments within cells and over different realizations of fading.

Figure 17 shows the number of nodes discovered versus the number of nodes actually deployed. Shown in dashed line is the reference case of no fading, with only a single transmission of every ECP. Solid and dotted lines correspond to a fading channel, with multiple ECP transmissions in two forms (with more or with less power, respectively).

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 four repetitions of the ECP 4, whereas 75 nodes were discovered using seven repetitions. When it comes to choosing the number of ECP repetitions, there appears to be an effect of diminishing returns, that is, there is more to be gained by going from 4 to 5 repetitions than from 5 to 7, and so on.

Figure 18 shows the percentage of times the discovery procedure was successfully completed (all nodes were discovered) for the same settings as those of Figure 17. For example, with 81 nodes, the discovery procedure was completed in about 75% of cases using seven ECP repetitions, whereas only in 38% of cases using four ECP repetitions. One could thus say that the chances of discovering all the 81 nodes are 75%, if seven ECP repetitions are used. Finally, we note that reliability of discovery decreases with the number of nodes. This is a consequence of the fact that fading is modeled as independent on different hops; hence, the chances of failure (premature ending at any given hop) increase with the number of hops. Our current simulation does not take into account spatial correlation of fading, which may affect the results.

Figure 19 shows the average total time that elapsed before the discovery procedure would end (possibly with not all the nodes discovered). Figures 20 and 21 show the corresponding energy consumption. These results are to be interpreted in conjunction with the number of nodes discovered (Figure 17). For example, with seven ECP repetitions made at the next higher power level, out of a total of 81 nodes, 75 were discovered in 38 min when the procedure ended having consumed 310 J. When the number of ECP repetitions was reduced to 4, only 60 nodes were discovered in 32 min using 200 J of energy. Observing the dotted curve corresponding to four ECP repetitions in Figures 19 and 17, we notice that it takes slightly less time to discover 30 nodes out of 81 then it takes to discover 29 nodes out of 64, which is explained by the fact that node density is higher (inter-node distance is shorter) in the former case.

5.4.2. Sensitivity / fading strength.

The question that remains to be addressed is that of protocol sensitivity to fading. We focus on two versions of the protocol, one with seven ECP repetitions and another with four ECP repetitions made at the next higher power

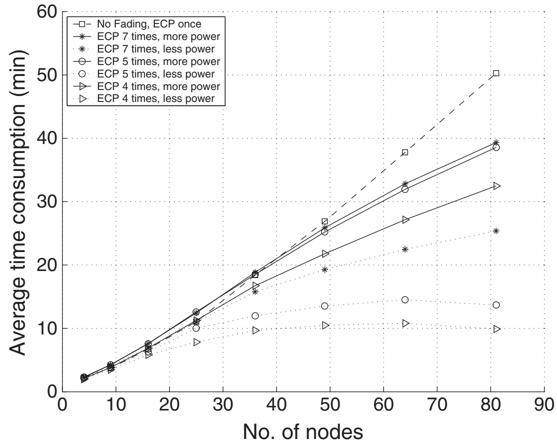

Figure 19. Average time consumption versus the number of nodes deployed.

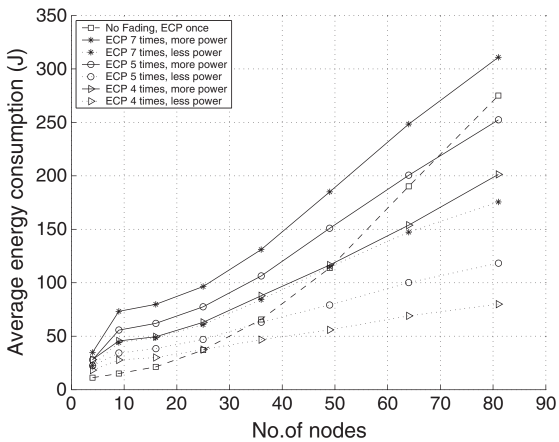

Figure 20. Average total energy consumption versus the number of nodes deployed.

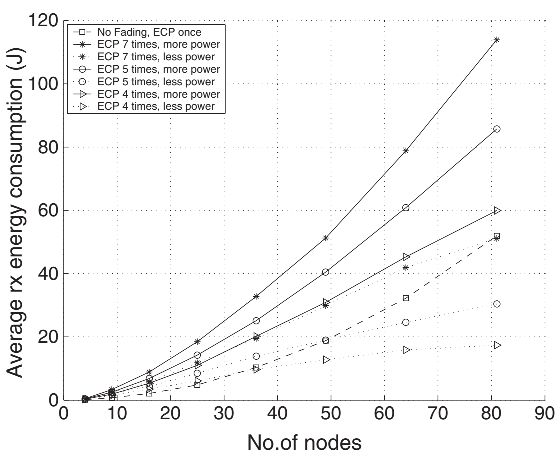

Figure 21. Average receiving energy consumption versus the number of nodes deployed.

level and investigate the performance under varying fading conditions described by the parameter A/σ.

Figure 22 shows the number of nodes discovered for different values of the ratio A/σ. As expected, the performance worsens (fewer nodes are discovered) as the fading strength increases, that is, as A/σ decreases. Transmitting multiple ECPs at the next higher power level remains effective in recovering the performance: with 81 nodes at A/σ=17, 66 nodes are discovered using seven ECP repetitions as compared with only 43 nodes using four ECP repetitions. Assuming the seven ECP versions, 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} but can further be overcome by investing more power.

Figure 23 shows the percentage of times the discovery procedure was successfully completed. We note that reliability is very sensitive to the fading strength, changing from about 50% to 75% as A/σ varies from 17 to 30.

Figure 24 shows the average time it took before the discovery procedure ended, whereas Figures 25 and 26 show the corresponding energy consumption. As before, these results have to be interpreted in conjunction with the number of nodes discovered (Figure 22). In other words, the fact that less time and energy 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 such as those of Figures 22, 24, and 25 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, energy, maximal power) and protocol adjustments (number of ECP repetitions or similar) needed to achieve a desired performance (reliability, time to complete).

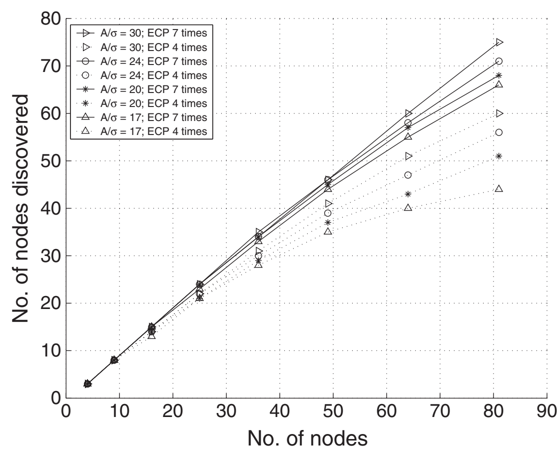

Figure 22. Number of nodes discovered versus the number of nodes deployed.

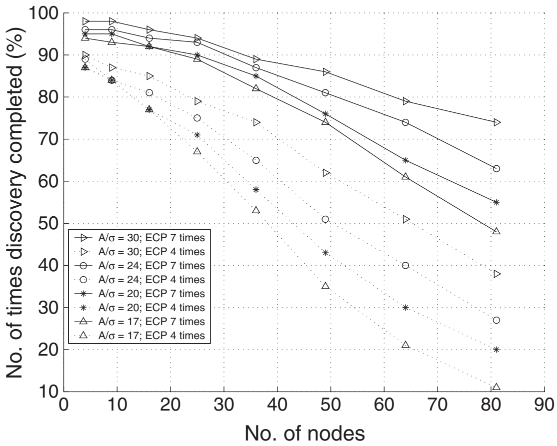

Figure 23. Percentage of times discovery procedure was successfully completed (all nodes discovered) versus the number of nodes in a fixed area.

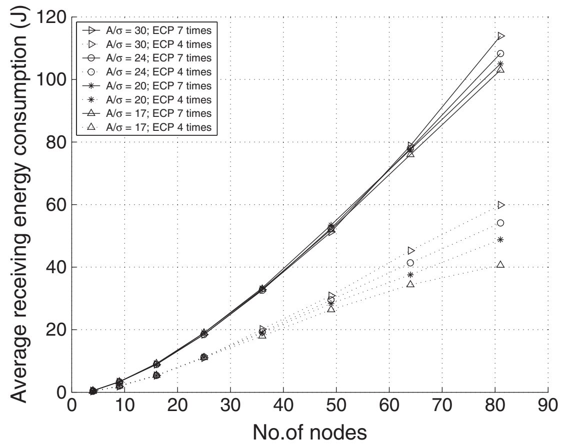

Figure 26. Average receive energy consumption versus the number of nodes.

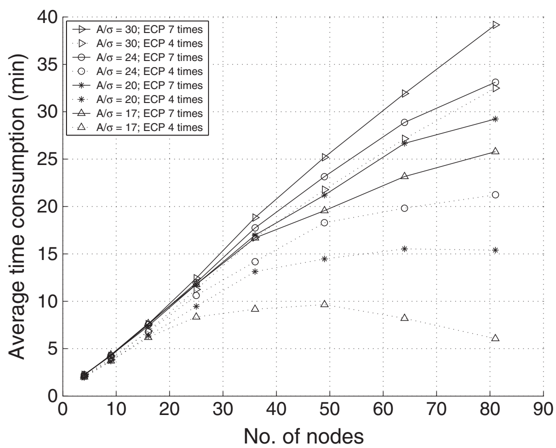

Figure 24. Average time consumption versus the number of nodes.

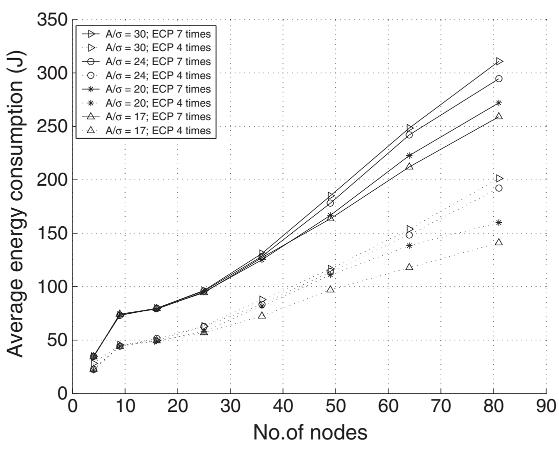

Figure 25. Average total energy consumption versus the number of nodes.

6. 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. The nodes operate in a distributed manner, that is, without a central station or a priori regulation of the channel access. There is a finite number of nodes covering a possibly large area that cannot be spanned in a single hop within the constraints of a finite power budget, that is, broadcast initialization is not an option.

Within this framework, a node discovery protocol was proposed in which the nodes use random access and transmit at minimum power level required to reach a particular neighbor. Conservation of energy is an important aspect in autonomous underwater acoustic systems, because the power is battery supplied, and transmission at greater-than-necessary power levels is not only wasteful but also creates interference which in turn increases the number of retransmissions.

Power control is implemented in discrete 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. Eliminating the need for scheduling is advantageous in underwater acoustic settings where location uncertainty, clock drifting, and long and variable propagation delays are present.

Network initialization begins arbitrarily with node 1 and 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 nodes have reached the maximum power level. Upon completion, each node has built a list of contacts, which includes at least the nearest neighbor, and regular network operation can commence. 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.

The performance of the protocol was assessed via simulation, by measuring the energy consumption and the time for completion. Comparisons with the benchmark case of broadcast initialization clearly demonstrate the benefits of the proposed scheme, whose energy consumption is much lower, whereas the time to completion, although increased, stays well within the limits needed to physically deploy a network of underwater nodes (robots, vehicles, floating sensors).

Underwater acoustic channels experience time-varying propagation conditions, whose impact on the protocol performance was assessed by including a fading model into the simulation analysis. According to our large-scale fading model, the channel between two nodes at a given time is characterized by the nominal attenuation (calculated for the given distance and frequency band) and an additional, log-normally distributed random component. The fading severity is quantified by the ratio of the nominal gain to the standard deviation of the random component.

Fading degrades the system performance, leading to packet loss that may cause a premature ending of the discovery procedure that leaves some nodes undiscovered. Simulation results quantify the system reliability by measuring the percentage of nodes discovered and the chances of complete discovery for varying fading conditions. Simple adjustments to the discovery protocol are proposed (packet repetitions) and 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.

ACKNOWLEDGEMENT

This work was supported in part by the NSF grants 0831728 and 0946610, the ONR grant N00014-09-0700, and the MIT Sea Grant College Program under NOAA grant NA060AR4170019/86. We would like to thank Dr. Hanu Singh of the Woods Hole oceanographic Institution for providing the experimental data used to generate the results of Figures 5 and 6 and Ms. Parastoo Qarabaqi for helping with the related data processing. We would also like to thank Mr. Josep Miguel Montana for the assistance with upgrading his simulation package [33].

REFERENCES

1. Farr N, Bowen A, Ware J, Pontbriand C, Tivey M. An integrated, underwater optical /acoustic communications system, In *OCEANS 2010 IEEE*, Sydney, NSW, 2010; 1–6.
2. Singh S, Webster SE, Freitag L, Whitcomb LL, Ball K, Bailey J, Taylor C. Acoustic communication performance of the WHOI micro-modem in sea trials of the Nereus vehicle to 11,000 m depth, In *OCEANS 2009, MTS/IEEE Biloxi - Marine Technology for Our Future: Global and Local Challenges*, Biloxi, MS, 2009; 1–6.
3. Green D. Acoustic modems, navigation aids, and networks for undersea operations, In *OCEANS 2010 IEEE*, Sydney, NSW, 2010; 1–6.
4. Freitag L, Grund M, Singh S, Partan J, Koski P, Ball K. The WHOI micro-modem: an acoustic communications and navigation system for multiple platforms, In *Proceedings of MTS/IEEE, OCEANS 2005*, Washington, DC, Volume 2, 2005; 1086–1092.
5. Singer A, Nelson J, Kozat S. Signal processing for underwater acoustic communications. *IEEE Communications Magazine* 2009; **47**(2): 90–96.
6. Pompili D, Akyildiz I. Overview of networking protocols for underwater wireless communications. *IEEE Communications Magazine* 2009; **47**(1): 97–102.
7. Heidemann J, Ye W, Wills J, Syed A, Li Y. Research challenges and applications for underwater sensor networking, In *Wireless Communications and Networking Conference (WCNC 2006)*, Las Vegas, NV, 2006; 228–235.
8. Chitre M, Shahabodeen S, Stojanovic M. Underwater acoustic communications and networking: recent advances and future challenges. *Marine Technology Society Journal, Marine Technology Society* 2008; **42**(1): 103–116.
9. Peleato B, Stojanovic M. Distance aware collision avoidance protocol for ad-hoc underwater acoustic sensor networks. *IEEE Communications Letters* 2007; **11**(12): 1025–1027.
10. Syed A, Ye W, Heidemann J. Comparison and evaluation of the T-Lohi MAC for underwater acoustic sensor networks. *IEEE Journal on Selected Areas in Communications* 2008; **26**(9): 1731–1743.
11. Chirdchoo N, Seng Soh W, Chua K. RIPT: A receiver-initiated reservation-based protocol for underwater acoustic networks. *IEEE Journal on Selected Areas in Communications* 2008; **26**(9): 1744–1753.
12. Zorzi M, Casari P, Baldo N, Harris A. Energy-efficient routing schemes for underwater acoustic networks. *IEEE Journal on Selected Areas in Communications* 2008; **26**(9): 1754–1766.
13. Pompili D, Melodia T, Akyildiz I. Distributed routing algorithms for underwater acoustic sensor networks.

IEEE Transactions on Wireless Communications 2010; **9**(9): 2934–2944.

14. Montana JM, Stojanovic M, Zorzi M. On joint frequency and power allocation in a cross-layer protocol for underwater acoustic networks. *IEEE Journal of Oceanic Engineering* 2010; **35**(4): 936–947.
15. Basagni S, Petrioli C, Petroccia R, Stojanovic M. Choosing the packet size in multi-hop underwater networks, In *OCEANS 2010 IEEE*, Sydney, NSW, 2010; 1–9.
16. Lucani D, Stojanovic M, Medard M. Random linear network coding for time division duplexing: when to stop talking and start listening, In *IEEE Infocom*, Rio de Janeiro, 2009; 1800–1808.
17. Zorzi F, Stojanovic M, Zorzi M. On the effects of node density and duty cycle on energy efficiency in underwater networks, In *OCEANS 2010 IEEE*, Sydney, NSW, 2010; 1–6.
18. Eustice RM, Singh H, Whitcomb LL. Synchronous-clock one-way-travel-time acoustic navigation for underwater vehicles. *Journal of Field Robotics, Special Issue on State of the Art in Maritime Autonomous Surface and Underwater Vehicles,* Chichester, UK, 2011; **28**(1): 121–136.
19. Rice J, Ong CW. A discovery process for initializing underwater acoustic networks, In *International Conference on Sensor Technologies and Applications*, Venice, 2011; 408–415.
20. Proakis J, Stojanovic M, Rice J. Design of a communication network for shallow water acoustic modems, In *Oceans Community Conference,* Volume 2*,* Baltimore, MD, 1998; 1150–1159.
21. Othman A, Adams A, Tsimenidis C. Node discovery protocol and localization for distributed underwater acoustic networks, In *Proceedings of the Advanced Int'l Conference on Telecommunications and Int'l Conference on Internet and Web Applications and Services (AICT-ICIW '06)*, Washington, DC, USA, 2006; 93.
22. Alonso G, Kranakis E, Wattenhofer R, Widmayer P. Probabilistic protocols for node discovery in ad-hoc, single broadcast channel networks, In *Proceedings of International Parallel and Distributed Processing Symposium (IPDPS'03)*, 2003.
23. Cornejo A, Viqarm J, Welch J. Reliable neighbor discovery for mobile ad hoc networks, In *Proceedings of the 6th International Workshop on Foundations of Mobile Computing (DIALM-POMC'10)*, 2010; 63–72.
24. Mc Glynn MJ, Borbash SA. Birthday protocols for low energy deployment and flexible neighbor discovery in ad hoc wireless networks, In *ACM international symposium on Mobile ad hoc networking & computing*, Long Beach, CA, USA, 2001; 137–145.
25. Liu D. Protecting neighbor discovery against node compromises in sensor networks, In *the 29th IEEE International Conference on Distributed Computing Systems (ICDCS '09)*, 2009; 579–588.
26. Vasudevan S, Kurose J, Towsley D. On neighbor discovery in wireless networks with directional antennas. *IEEE Infocom* 2005.
27. Berkhovskikh L, Lysanov Y. *Fundamentals of Ocean Acoustics*. Springer-Verlag: New York, NY, 1991.
28. Stojanovic M. On the relationship between capacity and distance in an underwater acoustic communication channel. *ACM Mobile Computing and Communications Review (MC2R), ACM,* 2007; **11**(4): 34–43.
29. Porter M. Bellhop code. Available online at http://oalib.hlsresearch.com/Rays/index.html.
30. Tomasi B, Casari P, Badia L, Zorzi M. A study of incremental redundancy hybrid ARQ over Markov channel models derived from experimental data. *International Workshop on Underwater Wireless Networks (WUWNet)* 2010; Article No. 4.
31. Yang WB, Yang TC. High-frequency channel characterization for M-ary frequency-shift-keying underwater acoustic communications. *Journal of the Acoustical Society of America* 2006; **120**: 2615–2626.
32. Kinsler LE, Frey AR, Coppens AB, Sanders JV. *Fundamentals of Acoustics*. Wiley: New York, 1991.
33. Montana JM. AUVNetSim: a simulator for underwater acoustic networks. *MIT Sea Grant Technical Report*, 2008. Available online at http://usersecegatechedu/jmjm3/publications/auvnetsimpdf.

AUTHORS' BIOGRAPHIES

Ashish Patil Ashish Prakash Patil was born in Mumbai, India, in 1986. He received his Bachelor of Engineering degree in Electronics and Telecommunication from the University of Mumbai, India, in 2008. He graduated with a Masters degree in Electrical Engineering from Northeastern University, Boston, MA, USA in 2010. His thesis was on ad hoc underwater acoustic networks. While at Northeastern University, he was awarded Research Assistantship by the Electrical and Computer Engineering Department. His research interests include ad hoc networks and wireless communications. Ashish is currently with Ericsson, Inc., where he works in the field of mobile communication systems.

Milica Stojanovic graduated from the University of Belgrade, Serbia, in 1988, and received the MS and PhD degrees in Electrical Engineering from Northeastern University, Boston, MA, USA in 1991 and 1993. After a number of years with the Massachusetts Institute of Technology, where she was a Principal Scientist, she joined the faculty of Electrical and Computer Engineering Department at Northeastern University in 2008. She is also a Guest Investigator at the Woods Hole Oceanographic Institution and a Visiting Scientist at MIT. Her research interests include digital communications theory, statistical signal processing and wireless networks, and their applications to underwater acoustic communication systems. Milica is an associate editor for the IEEE Journal of Oceanic Engineering and the IEEE Transactions on Signal processing.