Underwater Sensor Networking: Research Challenges and Potential Applications*

USC/ISI Technical Report ISI-TR-2005-603

John Heidemann Yuan Li Affan Syed
Jack Wills Wei Ye
USC/Information Sciences Institute

Abstract

This report summarizes our research directions in underwater sensor networks. We highlight potential applications to off-shore oilfields for seismic monitoring, equipment monitoring, and underwater robotics. We identify research directions in short-range acoustic communications, MAC, time synchronization, and localization protocols for high-latency acoustic networks, long-duration network sleeping, and application-level data scheduling.

1 Introduction

Sensor networks have the promise of revolutionizing many areas of science, industry, and government with their ability to bring computation and sensing into the physical world. The ability to have small devices physically distributed near the objects being sensed brings new opportunities to observe and act on the world, for example with micro-habitat monitoring [9, 31], structural monitoring [56], and wide-area environmental systems [47]. Industrial applications such as oil fields and production lines use extensive instrumentation, today often as carefully engineered SCADA systems, but increasingly with more rapidly deployed sensor networks [39]. Advances in reducing sensor cost and size imply that they can be inexpensive and small enough to be pervasive. The fact that these devices can communicate means that they can cooperate and relay data to remote users, operating unattended. Advances in energy efficiency mean that devices can observe long-term trends in their subjects.

While sensor-net systems are beginning to be fielded in applications today on the ground, *underwater* operations remain quite limited by comparison. Remotely controlled submersibles are often employed, but as large, active and managed devices, their deployment is inherently temporary. Some wide-area data collect efforts have been undertaken, but at quite coarse granularities (hundreds of sensors to cover the globe) [48]. Even when regional approaches are considered, they are often wired [16].

The key benefits of terrestrial sensor networks stem from wireless operation, self-configuration, and maximizing the utility of any energy consumed. We are currently exploring how to extend these benefits to *underwater sensor networks with acoustic communications*. It is instructive to compare current terrestrial sensor network practices to current underwater approaches. Terrestrial networks emphasize low cost nodes (around US$100), dense deployments (at most a few 100m apart), multihop communication, short-range communication; by comparison, typical underwater wireless communication today are typically expensive (US$10k or more), sparsely deployed (a few nodes, placed kilometers apart), typically communicating directly to a "base-station" over long ranges rather than with each other. We seek to reverse each of these design points, developing underwater sensor nodes that can be inexpensive, densely deployed, and communicating peer-to-peer.

Underwater sensor networks have many potential applications, including seismic monitoring, equipment monitoring and leak detection, and support for swarms underwater robots (explored in more detail in Section 3). Here we briefly consider seismic imaging of undersea oilfields as a representative application. One major reason to choose this application is that underwater sensor network is able to provide significant economic benefits over traditional technology. Today, most seismic imaging tasks for offshore oil fields are carried out by a ship that tows a large array of hydrophones on the surface [30]. The cost of such technology is very high, and the seismic survey can only be carried out rarely, for example, once every 2–3 years. In comparison, sensor network nodes have very low cost, and can be permanently deployed on the sea floor. Such a system enables frequent seismic imaging of reservoir (e.g. once every 3 months), and helps to improve resource recovery and oil productivity.

To realize these applications, an underwater sensor net-

---

[*] Authors are listed in alphabetical order.

3em This research is partially supported by the National Science Foundation (NSF) under grant number NeTS-NOSS-0435517, "Sensor Networks for Undersea Seismic Experimentation", by a hardware donation from Intel Corporation, and by CiSoft (Center for Interactive Smart Oilfield Technologies), a Center of Research Excellence and Academic Training and a joint venture between the University of Southern California and Chevron.

work must provide many of the tools that have been developed for terrestrial sensor networks: wireless communication, low-power hardware, energy conserving network protocols, time synchronization and localization, and programming abstractions. We can borrow many of these tools from ongoing, ground-based sensornet research. However, some of the challenges are fundamentally different. First, radio is not generally suitable for underwater usage because of extremely limited propagation (current mote radios transmit 50–100cm). While acoustic telemetry promises an alternative method of underwater wireless communication, off-the-shelf acoustic modems are not suitable for large-scale underwater sensor-nets: their power draws, ranges, and price points are all designed for sparse, long-range, expensive systems rather than small, dense, and cheap sensor-nets. Second, the shift from RF to acoustics changes the physics of communication from the speed of light (3×10^8m/s) to the speed of sound (around 1.5×10^3m/s)—a difference of five orders of magnitude. While propagation delay is negligible for short-range RF, it is a central fact of underwater wireless. This has profound implications on ranging and time synchronization. Finally, energy conservation of underwater sensor-nets will be different than on-ground because the sensors will be larger, and because some important applications require large amounts of data, but very infrequently (once per week or less).

We are therefore investigating three areas: *hardware*, acoustic communication with sensor nodes (Section 4); *protocols*, underwater-network network self-configuration, MAC protocol design, time synchronization, and ranging (Section 5); and *mostly-off operation*, data caching and forwarding and energy-aware system design and ultra-low duty cycle operation (also in Section 5). We believe that low-cost, energy conserving acoustic modems are possible, and that our focus on short-range communication can avoid many of the challenges of long-range transfer. Development of multi-access, delay-tolerant protocols are essential to accomplish dense networks. Low-duty cycle operation and integration and involvement of the application can cope with limited bandwidth and high latency.

Solving these constraints in the abstract is an underspecified problem; many solutions are possible, only some of which are likely useful. We begin by reviewing our overall architecture (Section 2) and the constraints placed on our work by several applications (Section 3).

2 System Architecture

Before describing specific applications, we next briefly review the general architecture we envision for an underwater sensor network. We begin by considering the rough capabilities of an individual underwater sensor node, how it interacts with its environment, other underwater nodes, and applications.

Figure 1 shows a logical diagram of a potential system. We see four different types of nodes in the system. At the lowest layer are the large number of sensor nodes to be deployed on or near the sea floor (shown as small yellow circles in the figure). They have moderate price, computing power, and storage capacity. They collect data through their sensors (e.g., seismic) and communicate with other nodes through short-range acoustic modems. They have batteries, but for long-term operation they spend most of their life asleep.

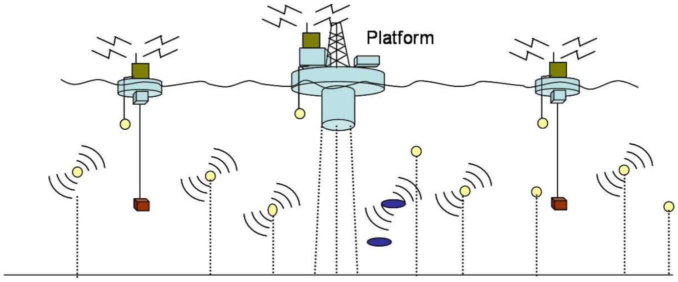

Figure 1. One possible approach to node deployment.

At the top layer are one or more control nodes with connections to the Internet, and possibly human operators. These control nodes may be positioned on an off-shore platform with power, or they may be on-shore; we expect these nodes to have a large storage capacity to buffer data, and access to ample electrical power. Control nodes will communicate with sensor nodes directly, via a relay node: a sensor node with underwater acoustic modems that is connected to the control node with a wired network.

In large networks, a third type of nodes, called *supernodes*, have access to higher speed networks. We are considering two possible implementations: first involves attaching regular nodes to tethered buoys that are equipped with high-speed radio communications to the base station, as shown in the figure. An alternative implementation would place these nodes on the sea floor and connect them to the base station with fiber optic cables. Regardless of the particular implementation, the important characteristic of supernodes is that they can relay data to the base station very efficiently. These supernodes allow a much richer network connectivity, creating multiple data collection points for the underwater acoustic network.

Finally, although robotic submersibles are not the focus of the current work, we see them interacting with our system via acoustic communications. In the figure, dark blue ovals represent multiple robots servicing the platform.

The computing power present in each node of a current sensor networks varies greatly, from 8-bit embedded processors, such as Berkeley Motes to 32-bit embedded processors about as powerful as typical PDAs, such as Intel Stargates to current 32- or 64-bit laptop computers. (Both motes and Stargates are currently available through Crossbow, Inc. at `xbow.com`.) We see Stargate-class computers as most appropriate for underwater sensor networks for several reasons. Their memory capacities (64MB RAM, 32MB flash storage) and computing power (a 400MHz XScale processor) is sufficient to store and process a significant amount of data temporarily, while their cost is moderate (currently US$600/each). Although Mote-class computers are attractive because extremely low cost and energy requirements, their very limited memory (4–8kB of RAM and 64–1024MB of flash storage) is a poor match for the requirements of underwater applications we describe in the next section (see

Section 3). Laptop computers have plenty of capability, but they cost twice or more what Stargates do, and add unnecessary keyboard and display.

Battery power and the ability to carefully monitor energy consumption is essential for the sensor node. It is essential that all components of the system operate at as low a duty cycle as possible; we expect to examine each layer of system software to minimize energy consumption, as described in Section 5. In addition to low duty cycle when operational, we expect to coordinate with the application to entirely shut off the node for very long periods of time, up to hours or days. We describe some of our plans in Section 5.5; we also expect to build on techniques such as those used by Intel [39].

We expect nodes to be deployed in several ways as shown in Figure 1. Where possible we expect nodes to be anchored to the ocean floor (the small yellow boxes in Figure 1). Tethers ensure that nodes are positioned roughly where expected (subject to drift or node damage), allowing optimization of placement for good sensor and communications coverage. We anticipate a tiered deployment (for example, as done in habitat monitoring [9]), where some nodes have greater resources. We expect the primary limiting resource to be communications capability, so our "supernodes" will have better communications, either by wired connections to each other and an external network, or by tethers to buoys with medium- or wide-area wireless communications (radios such as 802.11, or possibly satellite connections). These are shown as larger green boxes in Figure 1. Finally we expect some nodes to be mobile, either free floating, or connected to submersible robots (for example, see [4]), shown as dark blue ovals in Figure 1.

In a harsh, underwater environment, we must anticipate that some nodes will be lost over long deployments. Possible risks include fishing trawlers, underwater life affecting cables or nodes, or failure of waterproofing. We therefore expect basic deployments to include some redundancy in communication and sensing, so that loss of an individual node will not have wider effects. We expect that applications can cope with some non-uniformity in data. In addition, we expect that we will be able to recover from multiple failures, either with mobile nodes, or with human deployment of replacements.

Although many nodes will be tethered to one location, we expect that nodes may move, either due to drift of the anchor, disturbance from external effects. In other cases, nodes will move autonomously. We expect nodes to be able to localize themselves to determine their locations, as discussed later (Section 5.3).

Communications between nodes is an important focus of our work, because we see a large gap between our target deployment and currently available commercial, long-range, high-power, point-to-point, acoustics communications. We discuss our approach to low-power, short-range acoustic communications in Section 4. Equally important (and also unaddressed by current underwater work) are the networking protocols that allow underwater nodes to self-configure and coordinate with each other; we discuss these protocol issues in Section 5.

Finally, we have some basic assumptions about the applications that match these design. First, application benefit from local processing and temporary data storage. Storage can be used to buffer data to manage low-speed communications, "time-shifting" data collection from retrieval. In some cases, nodes benefit from pairwise communications and computation. Those capabilities are important for much of the infrastructure we propose, including time synchronization, routing, and fault-recovery. Finally, in most sensing applications, we expect the most data to be eventually relayed to the user via one or more links to the Internet or a dedicated network.

3 Applications

We see our approaches as applicable to a number of applications, including seismic monitoring, equipment monitoring and leak detection, and support for swarms underwater robots. We review the different characteristics of each of these below.

Seismic monitoring: A promising application for underwater sensor networks is seismic monitoring for oil extraction from underwater fields. Frequent seismic monitoring is of importance in oil extraction; studies of variation in the reservoir over time are called "4-D seismic" and are useful for judging field performance and motivating intervention.

Terrestrial oil fields can be frequently monitored, with fields typically being surveyed annually, or quarterly in some fields, and even daily or "continuously" in some gas storage facilities and permanently instrumented fields. However, monitoring of underwater oil fields is much more challenging, partly because seismic sensors are not currently permanently deployed in underwater fields. Instead, seismic monitoring of underwater fields typically involves a ship with a towed sonar array as the sensor and an air cannon as the actuator. Because such a study involves both large capital and operational costs (due to the ship and the crew), current underwater fields are evaluated rarely, typically every 2–3 years. As a result, interventions and asset management approaches suitable for terrestrial fields cannot easily be applied to underwater fields.

Use of a sensor network raises a number of research challenges: extraction of data, reliably, from distributed sensor nodes; localization, where each node to determines its location when it is deployed or should it move; distributed clock synchronization clocks for accurate data reporting; energy management approaches to extend sensor network lifetime for a multi-year deployment. We plan to address these challenges through low-power acoustic communication (Section 4) and new protocols for high-latency time synchronization, multiple access, scheduled data access, and mostly-off operation (Section 5). To understand the typical requirements of seismic sensing, we carried out a preliminary analysis of the data generated by seismic monitoring. Each sensor collects 3 or 4 channels of seismic data, each collecting 24 bits/channel at 500Hz. After a seismic event is triggered, we need to capture 8–10s of data. This leads to about 60kB of data per sensor per event. At our expected 5kb/s transfer rate, that implies about 120s/sensor to transfer this data over one hop.

Typical oilfields cover areas of 8km × 8km or less, and 4-D seismic requires sensor placement approximating a 50–100m grid. (We assume that current seismic algorithms can accommodate minor variations in sensor placement, provided they are known.) This implies a fairly large sensor network of several thousand sensors will be required to provide complete coverage. It also implies that a tiered communications network is required, where some supernodes will be connected to users via non-acoustic communications channels. Two possible implementations are buoys with high-speed RF-based communications, or wired connections to some sensor nodes. For a grid deployment we assume one supernode per 25 nodes (a 5x5 segment of the network), suggested all nodes are within two hops of a supernode and time to retrieve all data is about one hour (assuming each supernode can download data in parallel). Of course, one can trade-off the number of supernodes against the time required to retrieve the data. (With supernodes covering areas 4 hops wide, there is only one access point per 81 nodes, but data retrieval time will be much longer due to increased contention at the access point.) We expect to refine our design as we learn more about the problem.

Equipment Monitoring and Control: Underwater equipment monitoring is a second example application. Ideally, underwater equipment will include monitoring support when it is deployed, possibly associated with tethered power and communications, thus our approaches are not necessary. However, *temporary* monitoring would benefit from low-power, wireless communication. Temporary monitoring is most useful when equipment is first deployed, to confirm successful deployment during initial operation, or when problems are detected. We are not considering node deployment and retrieval at this time, but possibilities include remote-operated or robotic vehicles or divers.

Short-term equipment monitoring shares many technical requirements of long-term seismic monitoring, including the need for wireless (acoustic) communication, automatic configuration into a multi-hop network, localization (and hence time synchronization), and energy efficient operation. The main difference is a shift from bursty but infrequent sensing in seismic networks, to steady, frequent sensing for equipment monitoring.

Once underwater equipment are connected with acoustic sensor networks, it becomes an easy task to remotely control and operate some equipment. Current remote operation relies on cables connecting to each piece of equipment. It has high cost in deployment and maintenance. In contrast, underwater acoustic networking is able to significantly reduce cost and provide much more flexibility.

Flocks of Underwater Robots: A third and very different application is supporting groups of underwater autonomous robots. Applications include coordinating adaptive sensing of chemical leaks or biological phenomena (for example, oil leaks or phytoplankton concentrations), and also equipment monitoring applications as described above.

Communication for coordinated action is essential when operating groups of robots on land. Underwater robots today are typically either fully autonomous but largely unable to communicate and coordinate with each other during operations, or tethered, and therefore able to communicate, but limited in deployment depth and maneuverability.

We expect communications between underwater robots to be low-rate information for telemetry, coordination, and planning. Data rates in our proposed system are not sufficient to support full-motion video and tele-operation, but we do expect to be able to support on-line delivery of commands and the ability to send back still frame images.

4 Hardware for Underwater Acoustic Communications

We have described why underwater *acoustic communications* is an important alternative to radio-frequency (RF) communications for these networks.

At the hardware level, underwater acoustic communications is much like RF communications in air, but with a few key differences. In both systems we transmit a tone or carrier. This carrier is modulated by the data that we are sending. Common modulation methods include changing the carrier amplitude (AM), the carrier frequency (FM), and the carrier phase (PM). Modulation can occur in in a continuous (analog) or a stepped (digital) fashion. The primary differences between these modulation techniques lies in the complexity of the receiver, the bandwidth required, and the minimum acceptable received signal-to-noise ratio. This last parameter is usually expressed as E_b/N_o or *energy per bit over noise spectral density* [36, 55]. As an example, binary frequency shift keying, requires about 14 dB E_b/N_o for a 1×10^{-6} BER.

The received E_b/N_o depends on a few basic factors: the transmitter power, the data rate being sent, the noise level at the receiver, and the signal attenuation between the transmitter and receiver. We review each of these constraints next.

Transmit Power

There is no fundamental limit to transmitter power, but it can have a major effect on the power budget for the system. For energy efficiency and to minimize interference with neighboring transmitters we wish to use the smallest possible transmitter power.

Data Rate

This is a tradeoff in the system design, based on available power, and channel bandwidth. Because acoustic communications are possible only over fairly limited bandwdiths, we expect a fairly low data rate by comparison to most radios. We see a rate of currently 5kb/s and perhaps up to 20kb/s. Fortunately these rates are within an order of magnitude of RF-based sensor networks.

In application such as robotic control, the ability to communicate *at all* (even at a low rate) is much more important than the ability to send very large amounts of data quickly. In addition, in Section 5.5 we describe how application-level techniques can be used to maximize the benefits of even limited communications rates.

Noise Level

Noise levels in the ocean have a critical effect on sonar performance, and have been studied extensively. Burdic [6] and Urick [53] are two standard references. We are interested in the frequency range between 200 Hz and 50 kHz (the *mid-frequency band*). In this frequency range the dominant noise source is wind acting on the sea surface. Knudsen [28] has shown a correlation between ambient noise and wind force or sea state. Ambient noise increases about 5dB as the wind strength doubles. Peak wind noise occurs around 500 Hz, and then decreases about -6dB per octave. At a frequency of 10,000 Hz the ambient noise spectral density is expected to range between 28 dB/Hz and 50 dB/Hz relative to 1 microPascal. This suggests the need for wide range control of transmitter power.

Signal Attenuation

Signal attenuation is due to a variety of factors. Both radio waves and acoustic waves experience $1/R^2$ attenuation due to spherical spreading. There is also absorptive losses caused by the transmission media. For RF transmission, atmospheric losses are quite small. Absorptive losses in underwater acoustics are significant, and very frequency dependent. At 12.5 kHz absorption it is 1 dB/km or less. At 70 kHz it can exceed 20 dB/km. This places a practical upper limit on our carrier frequency at about 100kHz.

There are additional loss mechanisms, mostly associated with scattering, refraction and reflections. Stojanovic provides a very good overview of the challenges here [49]. A major difference between RF and acoustic propagation is the velocity of propagation. Radio waves travel at the speed of light, and, at ranges and frequencies typical for sensor networks, in a straight line. Acoustic transmission in water occurs at the speed of sound, which is around 1500 meters/sec. However the speed of sound in water varies significantly with temperature, density and salinity causing acoustic waves to travel on curved paths. This can create silent zones where the transmitter is inaudible. There can also be losses caused by multipath reflections from the surface, obstacles, the bottom, and temperature variations in the water and scattering from reflections off a potentially rough ocean surface.

Proposed Acoustic Communications Design

Many of these forms of loss are unique to acoustic communications at *longer* distances. In particular, multipath reflections, temperature variation, and surface scattering are all exaggerated by distance. Inspired by the benefits of short range RF communication in sensor networks, we seek to exploit *short-range underwater acoustics* where our only significant losses are spreading and absorption. We are developing a multi-hop acoustic network targeting communication distances of 50-500 meters and communication rates of around 5kb/s.

Using a simple FSK signaling scheme we anticipate sending 5kb/s over a range of 500m using a 30 mW transmitter output. The primary limitation is set be spreading loss and the background noise of the ocean. As with RF, we expect a combination of software and hardware techniques such as duty cycling can result in energy requirements a fraction of the basic transmit costs.

The use of low-power listening circuits has proven essential for RF-based sensor networks [43, 25, 17]. We are also developing a very low power *wakeup receiver* for our acoustic communications This receiver is not intended for data exchange, but only to detect acoustic energy in our channel and then signal our node that some is attempting to communicate. At this point we can turn on our data receiver/processor and enable communications. Our current hardware design using a dual gate FET configured as a cascode amplifier, with a passive filter and detector. The filter has a Q of 30, a center frequency of 18kHz. The circuit consumes 100 microamps at 5 volts (500 microwatts).

5 Protocols for High-Latency Networks

Acoustic communication puts new constraints *networks* of underwater sensor nodes for several reasons. First, the large propagation delay may break or significantly degrade the performance of many current protocols. The speed of sound in sea water is roughly 1.5×10^3m/s. The propagation delay for two nodes at 100m distance is therefore about 67ms. Second, the bandwidth of an acoustic channel is much lower than that of a radio. Efficient bandwidth utilization becomes an important issue. These constraints force us to review existing networking protocols and, in some cases, replace them with improved protocols designed explicitly for this high-latency environment. Finally, terrestrial networks can take advantage of rich existing infrastructure such as GPS and satellite communications networks. This constraint forces underwater sensor networks to be self-supporting in ways that terrestrial networks may not be. We next examine several research directions to provide this support for the underwater sensor networks.

5.1 Latency-Tolerant MAC Protocols

MAC protocols suitable for sensor networks can be broadly classified into two categories: scheduled protocols and contention-based protocols [61]. TDMA is a typical example of the scheduled protocols. It has good energy efficiency, but it requires strict time synchronization and is not flexible to changes in the number of nodes. Contention-based protocols are normally based on CSMA, and some collision avoidance mechanisms, such as RTS/CTS exchange, are also commonly used. Contention-based protocols have good scalability and adaptivity to changes in the number of nodes. Their energy efficiency can be improved by enabling low-duty-cycle operations on nodes, such as S-MAC [62, 63], STEM [44, 43], low-power listening [25], and asynchronous wake-ups [52, 65].

Currently the contention-based protocols with low duty cycles are widely studied by the sensor network community and results are promising. Although they are not optimized for ultra-low duty cycles as described in Section 5, they are still of great interest to many other applications. However, the large propagation delay in acoustic communications is particularly harmful to contention-based protocols for several reasons. First, it may take very long time for a node to detect its neighbor's transmission with its carrier sense. For example, suppose two neighboring nodes have a distance of 100m. If they try to send at about the same time, *e.g.*, triggered by the same sensing events, they need to listen for at least 67ms to avoid collisions. Furthermore, if a sender and a receiver exchange RTS and CTS, the overall propagation delay is tripled.

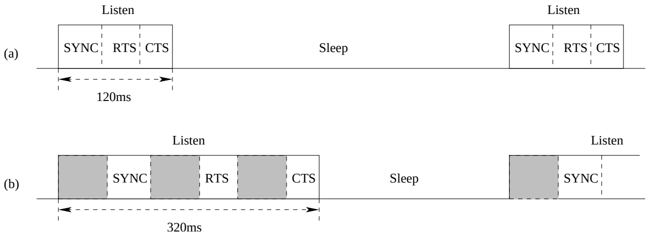

Figure 2. Modified S-MAC schedules to accommodate large propagation delay. (a) shows the listen window length currently implemented in TinyOS. (b) shows increased listen window to accommodate propagation delay of each packet.

Figure 2 shows the periodic listen and sleep schedule of a sensor node running S-MAC in low duty cycles. The top part (a) shows the length of the listen window in current implementation in TinyOS, which is about 120ms for listening SYNC, RTS and CTS packets. The bottom part (b) shows a naive extension to S-MAC where we modify the listening window to accommodate the propagation delays for each packet, now about 320ms. With this naive approach, a propagation delay will significantly increase the actual duty cycles of nodes, increase latency and decrease throughput, especially in multi-hop networks.

Clearly a major focus of MAC research will be to redesign media access protocols from the ground up to consider large propagation delays, rather than to simply adapt existing MAC protocols. First, we will examine the details of how the propagation delay affects energy efficiency, latency and throughput on existing protocols. Then, based on our understanding of the problem, we will develop new approaches to better accommodate the large propagation given the constraints in underwater sensor networks. Possible directions include designing new sleep and wake-up schemes, reducing control packet exchange, and combining contention-based transmissions with scheduled transmissions.

5.2 Time Synchronization

Time synchronization provides fundamental support for many protocols and applications. Without GPS, time synchronization algorithms have to be completely distributed over peering nodes. Several algorithms have been developed for radio-based sensor networks, achieving the accuracy of tens of microseconds [19, 22]. However, they assume nearly instantaneous wireless communication between sensor nodes, which is valid enough (0.33μs for nodes over 100m)for current RF-based networks.

In underwater sensor networks, the large propagation delay (for example, 67ms over 100m distances) becomes a dominant source of error for fine-grained time synchronization. Schemes like RBS [19] are built with the assumption of simultaneous reception of reference broadcasts, which results in synchronization error proportional to the propagation delay. TPSN [22] is not applicable since it fails to consider the effect of clock's skewing during the message exchange interval. Hence we have designed a time synchronization protocol, Time Synchronization for High Latency (TSHL), that can manage the high propagation latency induced errors while remaining energy efficient [51].

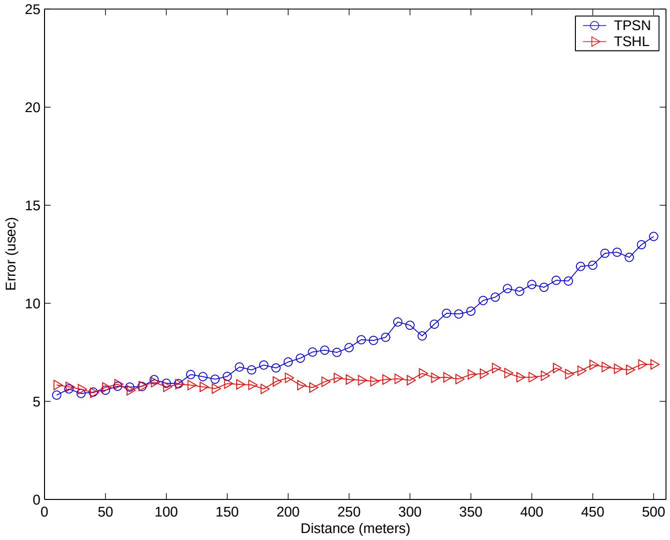

Figure 3. Comparison of clock synchronization error between TSHL and a TPSN-like protocol, immediately after a message exchange as distance between nodes increases.

The key idea in TSHL is that it splits time synchronization into two phases. In the first phase, nodes model their clock skew to a centralized timebase, after which they become *skew synchronized*. In the second phase they swap *skew compensated* synchronization messages to determine their exact offset. The first phase is impervious to the propagation latency, while the second phase explicitly handles propagation delay induced errors. This results in fast relative synchronization (end of phase 1), and also allows us to do *post-facto* synchronization. Both of these properties are highly desirable in our intended application (see 3.

We have evaluated TSHL in simulation to consider the effect of distances (and hence propagation latency), tolerance to clock skew, and design parameters of TSHL such as number of beacon messages used to estimate skew. At all distances, clock synchronization accuracy of TSHL is much better than RBS (by a factor of two or more), since RBS does not consider propagation latency at all. Figure 3 compares TSHL against TPSN, a protocol that does consider propagation delay. At short distances of less than 50m, synchronization accuracy of TSHL and TPSN are comparable, since for these distances clock skew during synchronization is minimal. At longer distances TPSN's failure to account for skew during synchronization causes increasing error in accuracy, up to twice the error at 500m. These values are immediately after the algorithm runs. Errors in clock estimation are magnified after synchronization, so TSHL is even more important when synchronization messages are done rarely to conserve energy.

We are in the process of implementing TSHL. We considered but rejected the alternative of substituting off-the-shelf, long-range acoustic modems to our intended short-range acoustic modems under development(Section 4. Primary reason being that such packages did not provide MAC layer timestamping capabilities; something that previous time synchronization schemes such as TPSN [22] demonstrated as essential for accurate synchronization. Instead, we have substituted in-the-air acoustic communication for underwater communication. We adopted the Cricket platform [35] due to its commercial availability and good support for low-level hardware access.

5.3 Localization

Localization is the process for each sensor node to locate its positions in the network. Localization algorithms developed for terrestrial sensor networks can be broadly divided into two classes. The first class is based on signal strength measurement [3, 5]. These algorithms are useful to give proximity information of nodes with low cost, but they are not able to provide accurate location information.

The second class is able to provide fine-grained location information, which is required by our seismic imaging application. These algorithms are based on measuring the propagation time of a signal, *i.e.*, the time-of-arrival (TOA) [42, 23]. Their basic principle for range measurement is the same as radar or sonar, but it is done in a distributed way among peering nodes. TOA measurement requires precise time synchronization between a sender and a receiver. Savvides et al [42, 23] use a radio signal to to synchronize the clocks of the sender and the receiver, *i.e.*, transmitting a radio signal at the same time a sound or ultrasound signal is transmitted. Since the radio propagation time is so small that the clocks of the two nodes are well synchronized; unfortunately, underwater networks will not be able to leverage this combination of RF and acoustic communication.

However, accurate range measurement can still be carried out if nodes have well synchronized clocks, and we will rely on our time synchronization work described in Section 5.2. Once the measurement is done among neighboring nodes, multilateration algorithms can be applied for each node to calculate its relative position to some reference nodes. The reference nodes can be the supernodes that are attached to buoys or off-shore platforms. If supernodes are placed on buoys, then they able to use GPS to obtain precise global locations, which can then be used as references to all underwater nodes. If supernodes are connected via wired networks, then we assume their locations can be surveyed when they are deployed and so they can again offer points of location reference.

While similar localization systems have been developed for terrestrial sensor networks [32], the accuracy of such systems need to be evaluated in the underwater environment. Unlike radio propagation, the speed of sound changes in the environment. Its actual value depends on temperature, pressure and salinity [12]. The propagation path may even be curved due to uneven temperature distribution. Moreover, node movement due to waves needs to be considered. All these factors affect localization accuracy and need to be studied.

5.4 Network Re-Configuration after Long Duration Sleeping

Undersea seismic monitoring of oil fields is an "all or nothing" application—periodically a seismic experiment will be triggered and all nodes must collect high-resolution seismic data for a few minutes, then a few months may go by with no activity. It would be extremely wasteful to keep the network fully operational for months at a time to support occasional measurements. Instead, we expect to put the whole network to sleep for the entire inactive period, reducing the duty cycle to a small percentage of deployment time (even before other optimizations are made). Similar approaches are also appropriate for long-term equipment monitoring, where nodes only need to check equipment status once a day or a week [39].

Prior work on energy conservation in sensor networks provides the illusion of constant access and 1–10% duty cycling via MAC-level sleep/wakeup [63] or low-power-listening [18, 34], and application-level approaches can get an additional factor of 2–4 from dense deployment [60, 11, 10, 44]. However, for the new class of applications, we must have a network and applications that can be completely shut down and quickly restarted, in effect, "sensor network suspend and resume". The major research issue is how to efficiently re-configure the network after a long sleep period.

Nodes agree on the same "resume" moment before their periodical long sleeps. But due to hardware limitations, they will wake up at different moments. When clock drift rate is 50ppm, the maximum clock drift without synchronization after 30 days is about 130 seconds. Thus nodes can wake up any time during 260 seconds period, so the network reconfiguration time after 30 days' inactivity is at least around 4 minutes!

There are two challenges in designing network reconfiguration protocols. Firstly, the re-configuration phases after long sleeps need to be as short as possible to restart the network quickly. Sensor nodes also need to stay energy efficient during these periods. And another challenge is to configure the network such that other protocols like MAC can resume quickly when the network resumes.

We propose two approaches. In the first approach, *low power listening with synchronized flooding*, right after nodes wake up asynchronously, they set up a timer twice the length of the maximum possible clock drift and perform low power listening (periodical short sampling) [25]. When the first node times out, all nodes should have restarted. It sends a network resume "UP" message immediately and the whole network starts flooding the message. Upon receiving the propagated message, nodes realize the network has resumed and data transmissions can begin without waiting for timeouts on their timers. In order to save energy during flooding, we synchronize nodes' sampling periods during flooding to reduce the overhearing overhead. In this approach, network can resume quickly by flooding and the network stays energy efficient since low power listening consume little energy.

Instead of flooding, in the second protocol we propose, *network configuration with requests and suppression*, the first node that is restart sets the network resume time, and the following nodes send requests to get the resume time from any already active nodes. In order to save energy, both request and replies are suppressed if possible—nodes listen for concurrent requests or replies and use them as their own.

The cost of reconfiguring a network must include the cost of brining up a fully functional MAC protocol. Asynchronous MAC protocols such as those based on low-power listening [18, 34] can start easily after the network resumes. We also wish to explore reconfiguration protocols that support MAC-level time synchronization, such as S-MAC. Our preliminary analysis suggests that we can achieve significant energy savings for both classes of MAC protocol compared to simply leaving nodes on idle listening during network reconfiguration.

And we are currently at the stage of implementing both protocols in TinyOS to verify their performance in real sensor networks.

5.5 Application-Level Data Scheduling

Besides energy constraints, acoustic networks also have very limited communications bandwidth. Today's off-the-shelf acoustic modems typically have the bandwidth between 5–20Kb/s. With applications like seismic imaging, all nodes will collect and try to send large amount of data that can easily overwhelm the network capacity. The research issue here is how to coordinate node's transmissions in an energy-efficient way that can best utilize the channel.

Current MAC protocols operating at 1–10% duty cycle provide the abstraction of a network that is always up by transparently delaying packets until the next awake period. This approach is not efficient for nodes to transmit large data at about the same time, as excessive MAC-level contention wastes bandwidth and energy. Instead we will explore explicit *application-level data caching and forwarding*. Building on the work of Delay Tolerant Networking [20], we plan to package sensor network readings and pass them from sensor node to sensor node.

While DTN outlines a generic architecture for store-and-forward data delivery, our seismic imaging application raises important application-level scheduling issues. For example, assume each sensor in Figure 4 must send 2.4MB of seismic data to the extraction node (indicated with an "X", assumed to have power and a surface network connection), and that each node can talk only to its immediate neighbors. Assuming an acoustic radio at 20kb/s and that each node as 2 minutes of 20kHz, 8-bit seismic data of data to send (ignoring overhead). Raw transfer time for one node is 16 minutes. Unscheduled transmission of all data would have all nodes competing to send and awake for at least 4 hours, a figure well in excess of a reasonable duty cycle, and in practice much longer due to channel contention at node X. If instead we schedule nodes to transfer data in the order given by node-id, than in the worst case, the nodes nearest X are each up for only 48 minutes (a savings of 77%), and edge nodes for only 16 minutes. Scheduling transmissions at the application level avoids excessive MAC-level contentions and can better utilize the channel and save energy.

6 Related Work

These research directions build on related work from several communities: the oil industry as a potential user of underwater sensor networks, oceanographic researchers who build underwater sensing and communication systems, and the wireless sensor network community. While summarizing existing work, we will also point out what is new in our proposed research.

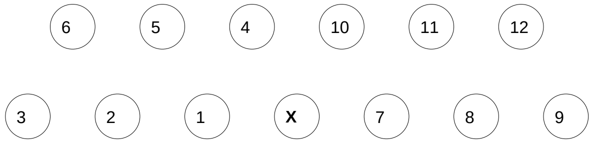

Figure 4. Extracting data from an underwater sensor grid.

6.1 Seismic imaging in oil industry

Three-dimensional (3-D) seismic imaging and monitoring is an important technology for oil exploration and reservoir management in the oil industry. Advanced reservoir management with 3-D seismic (sometimes 4-D with time series) can significantly improve resource recovery and oil productivity.

Today, most seismic imaging tasks for offshore oil fields are carried out by a ship that tows a large array of hydrophones on the surface [30]. A compressed-air gun generates a shock wave in the water. The wave travels down the sea floor and is reflected by different layers of the rock. The seismic signal is eventually received by each hydrophone on surface, and the data are processed coherently to form an image. Due to the high cost of such seismic imaging, it is only performed rarely, for example, once a year. An alternative way that has been used for underwater seismic is to deploy sensors underwater, which are connected by cables or fiber-optics [1]. The approach has the advantage of frequent data collection. However, it is very costly to lay cables underwater for a large area.

We propose a different approach for underwater seismic—using underwater wireless sensor networks. The sensor network consists of large numbers of smart sensors, and each of them has an embedded processor, sensors (seismic and others), storage memory, and acoustic communication devices. These nodes are battery powered, and are deployed in an ad hoc way without careful planning. Once deployed, the nodes will organize themselves into a multi-hop communication network, and gradually move sensing data back to users.

Our approach is new for undersea seismic imaging, and it has several advantages over existing ones. First, it is cost effective. These smart sensors are very cheap, so a large number of them can be deployed to cover a large monitoring area with enough density. Second, it is easy to deploy. It does not require special planning or extensive cable connections. Finally, it enables frequent seismic survey once the network is deployed.

6.2 Oceanographic research

Another related community is the oceanography, where researchers have developed underwater sensing and communication systems. An example is the Ocean Seismic Network program [48]. It developed seismic observatories in the deep ocean, as part of the Global Seismic Network (GSN). GSN has 128 observatories "uniformly" distributed on continents, islands or in the ocean, with a separation distance of 2000km. Its goal is to monitor a huge area on earth. In contrast, our sensor network covers a much smaller area, and nodes are densely deployed in an ad hoc fashion. In GSN, there is no direct communications between the sensing stations. They all directly send their data back to a central place. In sensor networks, the nodes will configure themselves to form a multi-hop communication network. In summary, the GSN is still the traditional way to do seismic imaging, but it covers a huge area including nodes in the ocean.

Underwater acoustic communication is another related area. The basic communication principles have been examined with acoustic channels in [37, 7, 49, 50]. Their major focus is the transmission range, bandwidth utilization and reliability with multi-path propagations. There are also experimental and commercial off-the-shelf acoustic modems available today, such as [40, 2, 29]. However, these modems are designed for long range communications (1–90km), and have weights of over 4kg. In our proposed hardware design, we will focus on short range, low-power modules in a small package. This capability is an enabling factor for long-lived sensor networks.

Networking protocols with acoustic communications are also studied in the literature. In [46], the authors reviewed MAC and routing protocols for wireless ad hoc networks. They also analyzed energy consumption with transmission range in acoustic transmissions, and pointed out that short-range, multi-hop relaying was the key for energy conservation. In [58], the authors studied the latency effects in acoustic communications and proposed a topology discovery algorithm for multi-hop communications. In [41], the authors proposed a clustering protocol with combined TDMA and CDMA for a group of autonomous underwater vehicles. These researches are based on an ad hoc networking model with small to moderate number of nodes. In contrast, our sensor network model consists of large numbers of nodes (hundreds to thousands), and our application has different requirements. The challenges we identified in this paper was not addressed by the existing work.

The NEPTUNE project [16] built an underwater sensor network with all nodes being connected by fiber-optic submarine cables. The cable supplied power to each node and provided high-speed communications. Follow-on work to the NEPTUNE network extended the wired network with some battery-powered nodes with acoustic communications [21]. In [21], the authors discussed the efficiency and reliability of modulations, and also briefly compared traditional MAC protocols, such as TDMA, FDMA and CDMA. The major difference of our sensor network model is that there will be no expensive cables laying on the sea floor. Most nodes will be cheap, small and battery-powered for easy deployment. Our work is focused on network self-organization, longevity, and multi-hop communications. The NEPTUNE network did not address the research problems we have identified.

6.3 Wireless sensor networks

Using wireless sensor networks for seismic imaging is not a new idea in the sensor network community. But all existing work are based on radio communications among sensors. Our goal is to extend sensor networking technology to underwater applications with acoustic communications.

So far, virtually all *platforms* developed for wireless sensor networks use radio communications. One of most widely used platforms is the UC Berkeley mote [26, 14], which is based on a 8-bit microcontroller, and a short-range, low-power radio. 32-bit platforms are normally embedded PCs, such as PC/104s and Stargates [15]. These big nodes do not have build-in radios, but can be connected with either motes or IEEE 802.11 cards. Although the radio propagation in water is very bad, the motes are still used by researchers in marine microorganism monitoring applications [8, 64]. We plan to extend sensor network platforms with a low-power, short-range acoustic communication device, so that large-scale underwater experiments and applications become possible.

There are several *networking protocols and algorithms* directly related to our proposed research. In fine-grained time synchronization algorithms, one approach is to synchronize different receivers to a common reference broadcast signal [19], and another one is based on sender and receiver pairs [22]. They are both designed for radio-based networks, where the propagation delay is negligible, but will break with acoustic communications. Fine-grained localization algorithms [42, 23] measure the TOA. Their success relies on fine-grained time synchronization, which is not available yet. Our approach will investigate the interactions of the fine-grained time synchronization and localization, and develop a combined algorithm.

Current research in the MAC layer is mainly on contention-based protocols, although TDMA protocols are also studied [38]. The major focus is energy efficiency, and several low-duty-cycle schemes have been proposed, such as S-MAC [62, 63], T-MAC [54], STEM [44, 43], low-power listening [25], and asynchronous wake-ups [52, 65]. The performance of existing MAC protocols needs to be re-evaluated with acoustic communications. New approaches need to be developed to accommodate large propagation delays.

Tiered architectures with heterogeneous nodes have also been studied. Example protocols include routing [57], clustering [24] and backbone formation [13, 59]. However, some problems are still not well investigated, such as the interactions between resources-limited nodes and resource-rich nodes in different layers. We plan to further investigate in this direction.

Prior work on *low-duty-cycle operation* aims to provide the illusion of constant network access. A common approach is the MAC-level sleep/wakeup [62, 63, 54, 44, 43, 25, 52, 65], which effectively enables duty cycles of 1–20%. An application-level approach exploits dense deployment by putting redundant nodes into sleep [60, 11, 10]. Now we are dealing with much longer sleep time with no application activities during sleeping. None of the above protocols are optimized for this type of applications. We must have new protocols to completely shut down and quickly restart the network.

An alternative to a network with low duty cycle operation is *no* duty cycle operation. Rather than build a connected network, nodes themselves may move and swap data using gossiping [27], or a designated node (a *data mule*) may move from node to node to explicitly gather data [45, 4, 33]. In situations where there is “free” movement, such as Zebra net [27] and DakNet [33] where nodes can hitch a ride on animals or vehicles, these approaches can be very inexpensive and in many ways ideal, assuming very high latencies (hours or days) can be tolerated. Deployment of actual robots is more challenging for many reasons: robot expense, continuing difficultly of autonomous navigation, and the high energy costs of actually physically moving a device. We see our approach at moderate-distance communication as providing an important complement to data mules. Furthermore, particularly in the case of underwater operation, data mules may benefit from acoustic communication rather than RF [4].

Another piece of related work is the Delay Tolerant Networking [20]. It outlines a generic architecture for store-and-forward data delivery. However, we need to further investigate important application-level scheduling issues in the underwater environment.

7 Conclusions

This paper has summarized our ongoing research in underwater sensor networks, including potential applications and research challenges.

Acknowledgments

We thank Chia-pin Hsiao, Ignatius Chan, and David Menda, all of Chevron, and Iraj Ershagi (USC Petroleum Engineering Dept.), for their input on the applications and application constraints. Gaurav Sukatme (USC Computer Science Dept.) and David Caron (USC Biology Dept.) provided feedback on biological applications and flocks of underwater robots.

8 References

[1] F. J. Barr. New seismic technique is improving hydrocarbon recovery. *Journal of Petroleum Engineering*, page 592, July 1996.

[2] Benthos, Inc., `http://www.benthos.com/pdf/Modems/ModemBrochure.pdf`. *Fast And Reliable Access To Undersea Data.*

[3] Jan Beutel. Geolocation in a picoradio environment. Master's thesis, UC Berkeley, July 1999.

[4] Vitaly Bokser, Carl Oberg, Gaurav S. Sukhatme, and Aristides A. Requicha. A small submarine robot for experiments in underwater sensor networks. In *Proceedings of the Fifth IFAC Symposium on Intelligent Autonomous Vehicles*, page to appear, Lisbon, Portugal, July 2004. International Federation of Automatic Control.

[5] Nirupama Bulusu, John Heidemann, and Deborah Estrin. Gps-less low cost outdoor localization for very small devices. *IEEE Personal Communications Magazine*, 7(5):28–34, October 2000.

[6] William S. Burdic. *Underwater Acoustic Systems Analysis*. Prentice-Hall, 1984.

[7] J. Catipovic, D. Brady, and S. Etchemendy. Development of underwater acoustic modems and networks. *Oceanography*, 6:112–119, March 1993.

[8] CENS. Cens research area: Monitoring of marine microorganisms. `http://www.cens.ucla.edu/portal/marine_microorganisms.html`, 2004.

[9] Alberto Cerpa, Jeremy Elson, Deborah Estrin, Lewis Girod, Michael Hamilton, and Jerry Zhao. Habitat monitoring: Application driver for wireless communications technology. In *Proceedings of the ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, San Jose, Costa Rica, April 2001. ACM.

[10] Alberto Cerpa and Deborah Estrin. ASCENT: Adaptive self-configuring sensor networks topologies. In *Proceedings of the IEEE Infocom*, page to appear, New York, NY, USA, June 2002. IEEE.

[11] Benjie Chen, Kyle Jamieson, Hari Balakrishnan, and Robert Morris. Span: an energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *Proceedings of the ACM International Conference on Mobile Computing and Networking*, Rome, Italy, July 2001. ACM.

[12] C-T. Chen and F.J. Millero. Speed of sound in seawater at high pressures. *Journal of the Acoustic Society America*, 62(5):1129–1135, 1977.

[13] W. Steven Conner, Jasmeet Chhabra, Mark Yarvis, and Lakshman Krishnamurthy. Experimental evaluation of synchronization and topology control for in-building sensor network applications. In *2nd ACM Int. Workshop Wireless Sensor Networks Applications (WSNA '03)*, San Diego, CA, September 2003.

[14] Crossbow Technologies, Inc., `http://www.xbow.com/`. *Mica2 Data Sheet.*

[15] Crossbow Technologies, Inc., `http://www.xbow.com/`. *Stargate Data Sheet.*

[16] John R. Delaney and Alan D. Chave. Neptune: A fiber-optic 'telescope' to inner space. *Oceanus*, 42(1):10–11, spring 2000.

[17] A. El-Hoiydi. Spatial TDMA and CSMA with preamble sampling for low power ad hoc wireless sensor networks. In *Proceedings of the SeventhInternational Symposium on Computers and Communications (ISCC)*, pages 685–692, July 2002.

[18] A. El-Hoiydi, J.-D. Decotignie, and J. Hernandez. Low power MAC protocols for infrastructure wireless sensor networks. In *Proceedings of the Fifth*, page xxx, Barcelona, Spain, February 2004. February.

[19] Jeremy Elson, Lewis Girod, and Deborah Estrin. Fine-grained network time synchronization using reference broadcast. In *Proceedings of the fifth USENIX Symposium on Operating System Design and Implementation (OSDI)*, Boston, MA, December 2002.

[20] Kevin Fall. A delay-tolerant network architecture for challenged internets. In *Proceedings of the ACM SIGCOMM Conference*, pages 27–34, Karlsruhe, Germany, August 2003. ACM.

[21] Lee Freitag, Milica Stojanovic, Matthew Grund, and Sandipa Singh. Acoustic communications for regional undersea observatories. In *Proceedings of the Oceanol-*

ogy International, London, U.K., March 2002.

[22] Saurabh Ganeriwal, Ram Kumar, and Mani B. Srivastava. Timing-sync protocol for sensor networks. In *Proceedings of the first ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 138–149, Los Angeles, CA, November 2003.

[23] Lewis Girod and Deborah Estrin. Robust range estimation using acoustic and multimodal sensing. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001)*, Maui, Hawaii, October 2001.

[24] Wendi Rabiner Heinzelman, Anantha Chandrakasan, and Hari Balakrishnan. Energy-efficient communication protocols for wireless microsensor networks. In *Proceedings of the Hawaii International Conference on Systems Sciences*, January 2000.

[25] Jason Hill and David Culler. Mica: a wireless platform for deeply embedded networks. *IEEE Micro*, 22(6):12–24, 2002.

[26] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for networked sensors. In *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 93–104, Cambridge, MA, USA, November 2000. ACM.

[27] Philo Juang, Hidekazu Oki, Yong Wang, Margaret Martonosi, LiShiuan Peh, and Daniel Rubenstein. Energy efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebranet. In *Proceedings of the Tenth International Conference on Architectural Support for Programming Languages and Operating Systems*, San Jose, CA, October 2002.

[28] V.O. Knudsen, R.S. Alford, and J.W. Emling. Digital communications. *Journal of Marine Research*, 7(12):410, 1948.

[29] LinkQuest, Inc., `http://www.link-quest.com/html/uwm_hr.pdf`. *Underwater Acoustic Modems*.

[30] W. Lynn. 3D, 4D, and beyond. *Journal of Petroleum Engineering*, page 35, January 1998.

[31] Alan Mainwaring, Joseph Polastre, Robert Szewczyk, and David Culler. Wireless sensor networks for habitat monitoring. In *Proceedings of the ACM Workshop on Sensor Networks and Applications*, pages 88–97, Atlanta, Georgia, USA, September 2002. ACM.

[32] W. Merrill, L. Girod, J. Elson, K. Sohrabi, F. Newberg, and W. Kaiser. Autonomous position location in distributed, embedded, wireless systems. In *IEEE CAS Workshop on Wireless Communications and Networking*, Pasadena, CA, September 2002.

[33] Alex (Sandy) Pentland, Richard Fletcher, and Amir Hasson. DakNet: Rethinking connectivity in developing nations. *IEEE Computer*, 37(1):78–83, January 2004.

[34] Joseph Polastre, Jason Hill, and David Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the 2nd ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 95–107, Baltimore, MD, USA, November 2004. ACM.

[35] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *Proceedings of the 6thACM International Conference on Mobile Computing and Networking*, pages 32–43, Boston, MA, USA, August 2000. ACM.

[36] John G. Proakis. *Digital Communications*. McGraw-Hill, 1995.

[37] A. Quazi and W. Konrad. Underwater acoustic communications. *IEEE Communications Magazine*, 20(2):24–30, March 1982.

[38] Venkatesh Rajendran, Katia Obraczka, and J.J. Garcia-Luna-Aceves. Energy-efficient, collision-free medium access control for wireless sensor networks. In *Proceedings of the first ACM Conference on Embedded Networked Sensor Systems (SenSys)*, Los Angeles, CA, November 2003.

[39] Nithya Ramanathan, Mark Yarvis, Jasmeet Chhabra, Nandakishore Kushalnagar, Lakshman Krishnamurthy, and Deborah Estrin. A stream-oriented power management protocol for low duty cycle sensor network applications. In *Proceedings of the IEEE Workshop on Embedded Networked Sensors*, pages 53–62, Sydney, Australia, May 2005. IEEE.

[40] Joseph A. Rice. Undersea networked acoustic communication and navigation for autonomous mine-countermeasure systems. In *Proceedings of the 5th International Symposium on Technology and the Mine Problem*, Monterey, CA, April 2002.

[41] Francisco Salva-Garau and Milica Stojanovic. Multi-cluster protocol for ad hoc mobile underwater acoustic networks. In *Proceedings of the IEEE OCEANS'03 Conference*, San Diego, CA, September 2003.

[42] Andreas Savvides, Chih-Chieh Han, and Mani B. Srivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the 7thACM International Conference on Mobile Computing and Networking*, pages 166–179, Rome, Italy, July 2001.

[43] Curt Schurgers, Vlasios Tsiatsis, Saurabh Ganeriwal, and Mani Srivastava. Optimizing sensor networks in the energy-latency-density space. *IEEE Transactions on Mobile Computing*, 1(1):70–80, 2002.

[44] Curt Schurgers, Vlasios Tsiatsis, and Mani B. Srivastava. STEM: Topology management for energy efficient sensor networks. In *Proceedings of the IEEE Aerospace Conference*, page to appear, Big Sky, Montana, USA, March 2002. IEEE.

[45] Rahul C. Shah, Sumit Roy, Sushant Jain, and Waylon Brunette. Data mules: Modeling a three-tier architecture for sparse sensor networks. Technical Report IRS-

TR-03-001, Intel Research—Seattle, January 2003.

[46] Ethem M. Sozer, Milica Stojanovic, and John G. Proakis. Underwater acoustic networks. *IEEE Journal of Oceanic Engineering*, 25(1):72–83, January 2000.

[47] David Steere, Antonio Baptista, Dylan McNamee, Calton Pu, and Jonathan Walpole. Research challenges in environmental observation and forecasting systems. In *Proceedings of the SixthACM International Conference on Mobile Computing and Networking*. ACM, Boston, MA, USA 2000.

[48] Ralph A. Stephen. Ocean seismic network seafloor observatories. *Oceanus*, 41(1):33–37, Spring 1998.

[49] M. Stojanovic. Recent advances in high-speed underwater acoustic communications. *IEEE Journal of Oceanic Engineering*, 21:125–136, April 1996.

[50] M. Stojanovic, J. G. Proakis, J. A. Rice, and M. D. Green. Spread-spectrum methods for underwater acoustic. In *Proceedings of the IEEE OCEANS'98 Conference*, volume 2, pages 650–654, Nice France, September 1998.

[51] Affan Syed and John Heidemann. Time synchronization for high latency acoustic networks. Technical Report ISI-TR-2005-602, USC/Information Sciences Institute, April 2005. Under submission.

[52] Yu-Chee Tseng, Chih-Shun Hsu, and Ten-Yueng Hsieh. Power-saving protocols for IEEE 802.11-based multi-hop ad hoc networks. In *Proceedings of the IEEE Infocom*, pages 200–209, New York, NY, June 2002. IEEE.

[53] R.J. Urick. *Principles of Underwater Sound for Engineers*. McGraw-Hill, 1967.

[54] Tijs van Dam and Koen Langendoen. An adaptive energy-efficient mac protocol for wireless sensor networks. In *Proceedings of the first ACM Conference on Embedded Networked Sensor Systems (SenSys)*, Los Angeles, CA, November 2003.

[55] Anthony D. Whalen. *Detection of Signals in Noise*. Academic Press, 1971.

[56] D. Whang, N. Xu, S. Rangwala, K. Chintalapudi, R. Govindan, and J. Wallace. Development of an embedded sensing system for structural health monitoring. In *Proceedings of the International Workshop on Smart Materials and Structures Technology*, January 2004.

[57] Andreas Willig, Rahul Shah, Jan Rabaey, and Adam Wolisz. Altruists in the PicoRadio sensor network. In *International Workshop on Factory Communication Systems*, Vasteras, Sweden, August 2002.

[58] Geoffrey G. Xie and John Gibson. A networking protocol for underwater acoustic networks. Technical Report TR-CS-00-02, Department of Computer Science, Naval Postgraduate School, December 2000.

[59] Ya Xu, Solomon Bien, Yutaka Mori, John Heidemann, and Deborah Estrin. Topology control protocols to conserve energy in wireless ad hoc networks. Technical Report 0006, Center for Embedded Networked Sensing, University of California, Los Angeles, January 2003.

[60] Ya Xu, John Heidemann, and Deborah Estrin. Geography-informed energy conservation for ad hoc routing. In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking*, pages 70–84, Rome, Italy, July 2001. USC/Information Sciences Institute, ACM.

[61] Wei Ye and John Heidemann. Medium access control in wireless sensor networks. In Taieb Znati, Krishna M. Sivalingam, and Cauligi Raghavendra, editors, *Wireless Sensor Networks*. Kluwer Academic Publishers, 2005. Also available as Technical Report, ISI-TR-580, USC Information Sciences Institute.

[62] Wei Ye, John Heidemann, and Deborah Estrin. An energy-efficient mac protocol for wireless sensor networks. In *Proceedings of the IEEE Infocom*, pages 1567–1576, New York, NY, June 2002.

[63] Wei Ye, John Heidemann, and Deborah Estrin. Medium access control with coordinated adaptive sleeping for wireless sensor networks. *ACM/IEEE Transactions on Networking*, 12(3):493–506, June 2004.

[64] Bin Zhang, Gaurav S. Sukhatme, and Aristides A. Requicha. Adaptive sampling for marine microorganism monitoring. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2004.

[65] Rong Zheng, Jennifer C. Hou, and Lui Sha. Asynchronous wakeup for ad hoc networks. In *Proceedings of the ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 35–45, Annapolis, Maryland, June 2003.