

# Time Synchronization for High Latency Acoustic Networks <sup>\*</sup>

ISI-TR-2005-602

Affan A. Syed  
Department of Computer Science  
University of Southern California  
asyed@isi.edu

John Heidemann  
Department of Computer Science  
University of Southern California  
johnh@isi.edu

## Abstract

Distributed time synchronization is an important part of a sensor network where sensing and actuation must be coordinated across multiple nodes. Several time synchronization protocols that maximize accuracy and energy conservation have been developed, including FTSP, TPSN, and RBS. All of these assume nearly instantaneous wireless communication between sensor nodes; each of them work well in today's RF-based sensor networks. We are just beginning to explore *underwater* sensor networks where communication is primarily via acoustic telemetry. With acoustic communication, where the propagation speed is nearly five orders of magnitude slower than RF, assumptions about rapid communication are incorrect and new approaches to time synchronization are required. We present Time Synchronization for High Latency (TSHL), designed assuming such high latency propagation. We show through analysis and simulation that it achieves precise time synchronization with minimal energy cost. Although at very short distances existing protocols are adequate, TSHL shows twice the accuracy at 500m, demonstrating the need to model both clock skew and propagation latency.

## 1 Introduction

Time synchronization is an important part of many distributed applications. In the classical networking and operating system literature it has been widely studied in context of

---

<sup>\*</sup>This research is partially supported by the National Science Foundation (NSF) under grant number NeTS-NOSS-0435517, "Sensor Networks for Undersea Seismic Experimentation", and by a hardware donation from Intel Corporation. It is also partly supported 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 ChevronTexaco.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

database queries [3] and security applications [17]. In the Internet, NTP is the canonical approach to provide distributed time synchronization [16].

Time synchronization is even more important in sensor networks, where applications such as acoustic beamforming [25] and target tracking [4] require collaboratively processing of time-sensitive data [7]. Sensor networks add the additional requirement that energy consumption of the synchronization protocol be minimized.

Many time synchronization protocols for sensor networks have been proposed recently [9, 7, 14, 23]. These mechanisms provide a high degree of precision while being reasonably energy efficient. The protocols adopt increasingly sophisticated approaches to reduce noise and account for latency in communications, but *all* assume that propagation latency is negligible and thus can be effectively factored out of design consideration.

While assumptions about propagation latency are appropriate for RF-based communications (where these protocols were intended), a number of researchers are beginning to explore *underwater* sensor networks (for example, see a recent survey [1]). Underwater, radio propagation is very limited (Mica-2 transmit range has been measured as less than 1m in fresh water [26]). *Underwater acoustic* communication provides a viable alternative [22]. A number of recent efforts have begun exploring acoustic communication for underwater sensor networks. In fact, commercial systems exist today, typically focused on long-range communication. We have begun examining *short-range* underwater acoustics (as described in Section 2.1).

Table 1 compares radio-based networks for general computing, sensor networks and satellites with short-range acoustic networks. While bitrates of underwater acoustics are comparable to current sensor network radios, propagation delay much higher. This delay is due to the five-orders-of-magnitude difference in the speed of sound in water compared to RF propagation: 1500 m/s compared to  $3 \times 10^8$  m/s. Underwater acoustic communication can be thought of as a "long slim pipe", with the worst features of satellite and RF-based wireless links.

The high propagation delay of underwater acoustics is especially hazardous for time synchronization algorithms. While NTP tolerates large delay [16], it does not consider energy consumption issues. Sensor-network-based time synchronization protocols consider energy consumption, but

Table 1. Comparison of Link Characteristics

Characteristic	Satellite	802.11 RF	Chipcon RF	underwater acoustic (short range)
Bit Rate	155 Mb/s	11 Mb/s	20–50kb/s	20–50 kb/s
Typical BER	$10^{-10}$	$10^{-5}$	$10^{-5}$	$10^{-2}$ [20]
Propagation Delay	~120 ms	< 1 $\mu$ s	< 1 $\mu$ s	~300 ms
Distance	~ 42,000km	<3km	< .5km	< .5km

all current protocols are optimized for RF-based networks, assuming nearly instantaneous and simultaneous reception (RBS [7], FTSP [14]) or ignore clock drift during synchronization (TPSN [9], LTS [23]).

The contribution of our work are to identify the constraints of short-range, underwater acoustic communication for sensor networks (Section 2), to quantify the inaccuracies this creates in current time synchronization protocols (Section 4), and to introduce a new protocol, *Time Synchronization for High Latency* (TSHL), that compensates for high-latency communication while minimizing energy consumption (Section 5).

The key idea in TSHL is that it splits time synchronization into two phases. In the first phase, nodes swap synchronization messages to determine a common time base. In the second base they estimate clock skew. Both phases consider propagation latency.

We demonstrate the importance of this approach by analysis (Section 4) and simulations (Section 6). Analysis and simulations show that RBS and FTSP are not applicable to high-latency networks because they do not consider propagation latency at all. Through simulations we compare TSHL to TPSN, the closest current protocol. An implementation of TSHL is underway using Cricket nodes. We discuss how this in-the-air acoustic network will approximate the latency of an underwater network in Section 7.

## 2 Background

### 2.1 Underwater Acoustic Sensor Networks

1

Although radio-frequency communication is ideal for surface sensor networks, typical RF propagates very poorly underwater. Experimental tests with 433MHz Chipcon radios show propagation of less than 1m at full Mica-2 transmit power [26]. Although long-wavelength RF can penetrate water, it requires large transmit power and large antennae, making it inappropriate for small, low-power sensor nodes. This limitation suggested the use of robotic data mules to Zhang et al. [26], but the navigation and control problems there remain open issues.

These constraints suggest *acoustic communication* as an important alternative to RF. Stojanovic provides a very good overview of the options and challenges here [21]. 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




Figure 1. Underwater Acoustic Sensor Network Topology

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 wave to travel on curved paths. That traditional approaches in this field have focused on medium- and long-range approaches considering 1-10km and 10-100km distances. Stojanovic identifies a number of sources of signal loss: spreading and absorption loss (a function of distance and frequency); multipath reflections from the surface, obstacles, the bottom, and temperature variations in the water; noise due to artificial and natural sources; and scattering from reflections off a potentially rough ocean surface.

Many of these forms of loss are unique to acoustic communication 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*. We are developing a multi-hop acoustic network targeting communication distances of 50–500m and communication rates of 5–20kb/s.

Using a simple FSK signaling scheme we anticipate sending 5 kb/s over a range of 500m, using a 30mW transmitter output. The primary limitation on performance is set by 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.

Figure 1 shows a conceptual view of a target system using underwater acoustic communication. A relatively dense deployment of sensor nodes (small yellow circles) are connected by wireless underwater acoustic links. These communicate, possibly via multiple hops, with each other and with tethered nodes (small green squares) at buoys to users on a platform or ship. Where possible we can exploit a hybrid network with high-speed RF communication on buoys (large green squares), or with underwater or surface-based robotic or manually operated data mules (not shown).

<sup>1</sup>The material in this section is joint work with Jack Wills, Wei Ye, and others in the SNUSE group at USC/ISI.

While our work in prototyping underwater acoustics is ongoing and this multi-hop, hybrid network clearly requires much more work, our acoustic modem design places bounds on the latency and bitrates that we can expect. With 50–500m distances and a 1500m/s speed-of-sound we expect 300ms worst-case propagation latencies (ignoring any MAC effects).

## 2.2 Need for Time Synchronization: Clock offset and skew

Two challenges face synchronization of distributed clocks. First, they must be synchronized to a single common event in absolute time or *offset* (shown as  $b$  in Figure 3). Second, because clocks are imperfect and run at slightly different rates, one must determine the *skew* of a given clock relative to some absolute frequency.

Since computers boot at different times, there must be some way for distributed computers to determine a common offset. Offset can be determined by a single message exchange, provided we can compensate for any sources of non-deterministic latency in the path (described in Section 2.3). As examples, NTP compensates for jitter in the path by using control theoretic mechanisms [16], while RBS uses a jointly visible broadcast packet to synchronize nodes [7].

Time in most modern, inexpensive computers is derived from oscillating frequency of a quartz crystal. Due to environmental variation (temperature, humidity, etc.) or minor manufacturing differences, variations in crystal oscillation frequency on the order of 15–25 parts per million are common [2]. Thus, even nodes that are synchronized to a common offset will drift out of synchronization over time. In Figure 3 the dotted  $y = x$  line represents an ideally ticking clock where time on the  $x$ -axis matches time on the  $y$ -axis perfectly. Two clocks ticking at different rates with ratio  $a$  will appear as a line with a different slope  $y = ax$ . The ratio  $a$  represents clock skew. Protocols handle clock skew by estimating and compensating for it. For example, NTP, keeps the long term average skew error low by using a phase-locked loop to correct the local clock frequency [16], while RBS explicitly calculates this skew using linear regression [7].

With reasonable quality oscillators, clock skew is generally small enough that it can be ignored over short time intervals. Thus existing protocols ignore skew when synchronizing offset. While appropriate for RF-based networks, we show that this assumption does not hold for slower acoustic communication (see Section 4). For these networks, one must consider skew even *during* synchronization.

## 2.3 Sources of Errors in Time Synchronization

The major cause of error in time synchronization schemes is the non-determinism in the latency estimates of the message delivery delay. A description of sources of variability was first described by Koeptz and Schwabl [12] and extended recently by Horauer et. al. [11] incorporating physical layer jitter that cannot be overlooked for high precision time synchronization. We briefly review these sources of error below and in Figure 2.

1. **Send Time:** The delay in the packet traversal from the message assembly at the application layer all the way




Figure 2. Sources of error in estimating message latency

down to MAC layer. Highly non-deterministic.

2. **Access Time:** Is the channel contention time, that in dense broadcast medium such as ours can be in the order of hundreds of milliseconds. Least deterministic part of the message delivery.
3. **Interrupt Handling Time:** The delay between the radio chip raising and the microcontroller responding to an interrupt. Can be an issue if interrupts are disabled on the microcontroller.
4. **Transmission and Reception Time:** The delay in sending or receiving the entire length of the packet over the channel. Largely deterministic, is a function of bandwidth and packet size.
5. **Propagation Time:** The delay, for a particular symbol of the message, in traversing all the way to the receiver. The propagation time can be deterministic if the speed of propagation is assumed constant, and endpoint location is known, or if synchronization exchange is performed with assumption of path symmetry. This delay can be significant in the AU channel since we cannot assume that clocks will not skew over packet exchanges.
6. **Encoding and Decoding Time:** The time taken by the radio chip to encode/decode and transform a part of the message to/from electromagnetic waves. This time is deterministic and is in the order of hundred microseconds [14].
7. **Byte Alignment Time:** The delay because of the different byte alignment at the receiver. This time is deterministic and can be computed on the receiver side from

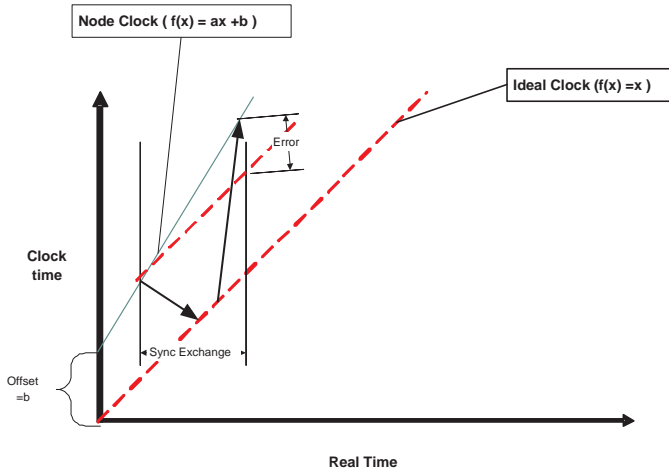
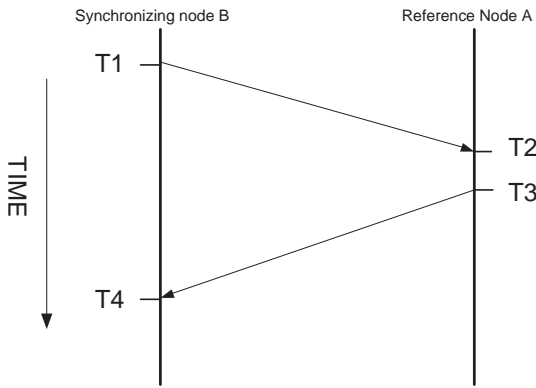


Figure 3. Effect of clock skew



$$\text{Phase offset} = [(T2-T1) - (T4-T3)]/2$$

$$\text{Propagation Delay} = [(T2-T1) + (T4-T3)]/2.$$

Figure 4. Sender-Receiver Synchronization

the bit offset and the speed of the radio.

8. **Receive Time:** Time for the incoming message to traverse up till the receiver application. Highly variable and varies for each (stack,OS) pair.

Existing time synchronization schemes (reviewed in the next section) focus on eliminating or accounting for these sources of error. Schemes typically differ due to differing assumptions in which sources of variation are dominant in different domains, and due to different approaches to eliminate the sources of error.

### 3 Related Work

An important notion in time is that it has to be relative to a given reference standard. Lamport clarified the relationship between computer events and global reference time [13]. We focus on time synchronization to a reference value motivated by the need to relate computer sensed events to the outside world.

At the most fundamental level, there are just two schemes to synchronize clocks: Sender-Receiver (Figure 4) and

Receiver-Receiver (Figure 5). All schemes operate within these two basic frameworks. In addition, some schemes synchronize against an external time reference, while others synchronize nodes to some arbitrary internal reference.

Network Time Protocol (NTP) is widely used in the Internet. It is distinguished by working well over paths with high latency and high variability [16]. The NTP protocol has a long-term, bi-directional exchange of time information to estimate both offset and skew. It incrementally adjusts the local clock frequency to align it with the reference time base. Unfortunately, NTP is a poor match for sensor networks for several reasons. First, it assumes communications are relatively inexpensive, while sensor networks are bandwidth and energy constrained. Second, it is designed for constant operation in the background at low rates. (At a maximum polling rate of 16 sec, NTP took around an hour to reduce error to about  $70\mu\text{s}$  [6]). By comparison, TSHL exchanges number of broadcast beacons to compute skew and then perform one bidirectional exchange to compute a skew-corrected offset. In some sense, TSHL and NTP possess the same information, however TSHL reduces energy consumption by replacing long-term bidirectional communication with a smaller number of unidirectional, broadcast beacons. In addition, TSHL is not constrained by portability requirements and so can exploit MAC-level timestamping as TPSN does.

An interesting extension of NTP considers the Interplanetary Internet (IPin) [15]. The protocol iNTP, as proposed [15], assumes very high latencies but very predictable node position and movement (for example, predictable trajectories of satellites). While we expect the approximate locations of underwater nodes to be known with some accuracy, we expect ocean currents and environmental effects to render position information insufficiently reliable.

An alternate Internet based protocol was clock skew compensation for streaming audio in the Internet [8]. Faced with large and varying path delays, Fober demonstrates how to model the drift of between node clocks without modeling the offset. He uses statistical measures to remove the high jitter expected for their application. Although we could apply these techniques in underwater acoustic networks to remove this high jitter, but they can also be removed considerably in our point to point network through MAC layer time stamping.

The research closest to our work is time synchronization effort in the sensor networks community. Underwater sensor networks share many of the design goals of surface sensor networks. Energy conservation and longevity given a fixed power budget are common goals.

Reference Broadcast Synchronization (RBS) introduced receiver-receiver synchronization, completely eliminating transmitter side uncertainties as described in Section 2.3 [7]. RBS accounts for clock skew by modeling the clock with linear regression on multiple *Reference Broadcasts*. Its accuracy is quite good, about  $6\mu\text{s}$  on IPAQ's with 802.11 cards. In addition, it introduces the concept of *post-facto synchronization*, allowing correction of clock errors after data collection rather than ahead of time. The RBS is not applicable to our environment because its central algorithm is built on the simultaneous reception of reference broadcasts at all nearby




Figure 5. Receiver-Receiver Synchronization

nodes. With underwater acoustics there are large variations in propagation time between nodes, resulting in synchronization error proportional to the propagation delay, up to 100ms or more.

Timing-sync Protocol for Sensor Networks (TPSN) exploited cross-layer optimizations, minimizing the sender and receiver side uncertainty by time stamping packets at the MAC [9]. Since TPSN uses a two-way exchange for synchronization, it can factor out most major sources of non-determinism (the Send, Receive and Access time) and propagation delay. The TPSN authors report achieving accuracy of  $8\mu\text{s}$  on Mica-2 motes. Like NTP, TPSN requires formation of a hierarchy, where each node synchronizes to its parents. However, it does not model skew of the local clock against the reference; instead it only computes offset. It therefore requires frequent resynchronization to correct for drift due to variation in clock skew, and it cannot do post-facto synchronization. The primary reason TPSN is not applicable to our environment is that it does not consider the effect of clock skew during message exchange. Although this effect is tiny during radio-based synchronization, it causes inaccuracies proportional to message propagation latency rises. We show in Section 6 that at moderate distances  $\sim 300\text{m}$  this error can be nearly 30% worse (see Figure 7).

Flooding Time Synchronization Protocol (FTSP) takes an additional step towards avoiding timestamp uncertainty by timestamping in the MAC and radio message layer in multiple places and using the average of these values to account for byte alignment jitter [14]. It uses a variation of sender-receiver synchronization. A sender broadcasts its global reference time to all receivers. Receivers use linear regression over multiple broadcasts to model their clock skew and offset. Compared to RBS, this approach reduces the number of message exchanges that need to take place for synchronization, since the message exchange is not between every pair of nodes. For underwater acoustic networks, however, FTSP

suffers from the same problem as RBS, in that it assumes near instantaneous message propagation. This delay is large and variable in underwater acoustic networks.

#### 4 Quantifying the Challenges of High Latency Links

Ideally we could simply use an existing time synchronization protocol in underwater acoustic networks. We next develop a simple analytic model to demonstrate why existing protocols such as RBS, TPSN, and FTSP do not work well for high-latency links. In this analysis we focus only on propagation latency as a source of error in time synchronization. Clearly this assumption is not correct in general, however, since existing schemes address other sources of error and are already accounted for by most existing protocols.

First we define a simple notation for realistic (inaccurate) clocks. For a node  $S$ , we model its uncorrected clock as  $f_S(t)$  and its corrected time as  $\hat{f}_S(t)$ :

$$\begin{aligned} f_S(t) &= a_S t + b_S \\ \hat{f}_S(t) &= f_S(t) + \beta_S(t) \end{aligned} \quad (1)$$

These are first order linear function of its skew  $a_S$  and offset  $b_S$ , where  $t$  is the global reference time, and  $\beta(t)$  is the correction factor calculated at  $t$ .

**Protocols based on one-way exchanges:** We first quantify the proportionality of error in both RBS and FTSP to the propagation delay by taking a simple case where a perfectly synchronized beacon node  $f_B(t) = \hat{f}_B(t) = t$  sends synchronizing pulses with its sending time to an uncorrected node  $R$ . Note that this is very similar to what happens in FTSP, and with slight modification can be applied to RBS as well.

With no propagation delay, the correction factor calculated by the synchronization protocol is:

$$\beta_R(t) = f_B(t) - f_R(t) = (1 - a_R)t - b_R$$

which will synchronize node  $R$ 's clock to the correct time, using Equation 1. The clock skew can be computed with multiple exchanges by observing how  $\beta_R(t)$  changes.

However, given a propagation delay of  $d$  between beacon and node  $S$ , the computed correction factor is:

$$\beta(t+d) = f_B(t) - f_S(t+d) = (1 - a_S)t - (a_S d + b_S)$$

Based on Equation 1 alone, the corrected time is offset incorrectly by  $d$ . To correct for this error we need to estimate propagation delay. To stay within the RBS or FTSP models we must estimate this delay without sending additional messages. A prior computation of underwater propagation speed is quite difficult. Node locations may be known, but are likely inexact due to placement error or node movement. Underwater node localization is an open problem, and often requires synchronized clocks for computation. Finally, knowledge of acoustic propagation speed assumes that the

nodes have capabilities to measure, at the very least, the temperature and pressure. For these reasons, a priori computation of delay seems quite challenging.

**Protocols based on two-way exchanges:** We next quantify the effect of the propagation delay on protocols that allow two-way exchanges, such as NTP and TPSN. The key element of these protocols are that they can factor out propagation delay via a two-way message exchange. However there is one assumption they make about sender's clock not skewing during the message exchange. Consider the sender  $S$  and the receiver  $B$  (perfect clock) exchanging timestamps as shown in Figure 4, with  $T_1 = f_S(t_1)$ ,  $T_2 = f_B(t_1 + d)$ ,  $T_3 = f_S(t_3)$ , and  $T_4 = f_S(t_3 + d)$ . With that scenario, the clock model for the sender  $S$  will be:

$$\begin{aligned} f_S(t) &= t + b_S \\ \beta_S(t) &= \frac{(T_2 - T_1) - (T_4 - T_3)}{2} \\ \hat{f}_S(t) &= f_S(t) + \beta_S(t) \end{aligned} \quad (2)$$

With no skew, and with our assumption of none of the sources of error present, this will perfectly synchronize the nodes.

We will now show that if we consider skew it does have an effect on the error as either the skew or the duration for the message exchanges is increased (intuitively you can grasp this by looking at the message exchange shown in Figure 3).

$$\begin{aligned} f_S(t) &= at + b_S \\ \beta(t_3 + d) &= \frac{(1 - a_S)(t_1 + t_3 + d) - 2b_S}{2} \end{aligned} \quad (3)$$

The delay compensating clock at  $S$  as:

$$\hat{f}_S(t_3 + d) = f_S(t_3 + d) + \beta(t_3 + d)$$

Using Equation 4 below shows that this leads to an Error =  $(t_3 + d) - \hat{f}_S(t_3 + d)$ , that is proportional to the skew and the packet exchange period.

$$Error = \frac{(1 - a_S)((t_3 - t_1) + d)}{2} \quad (4)$$

For Berkeley notes, the upper bound given in the datasheet [2] is 40ppm i.e. a clock in mote can loose up to  $40\mu s$  in a second. This can translate to a drift of around 15-20 $\mu s$  when the nodes are  $\sim 400m$  apart. This significantly affects the accuracy of time synchronization, as predicted by Equation 4, and our simulation will show the degradation of TPSN precision with increasing distance, or skew, between nodes.

## 5 Design of Time Synchronization for High Latency Channels(TSHL)

While prior time synchronization protocols addressed many sources of error in estimation, only NTP faced large propagation delay, and it is not appropriate for sensor networks (see Section 3). We now present an alternate time syn-

chronization algorithm for sensor networks that can manage high propagation delays while remaining energy efficient.

### 5.1 Overview and Assumptions

TSHL is a two phase protocol. The the core idea is to *first* model the skew of a node's clock so that each node is *skew synchronized*. We compute skew by performing linear regression over multiple beacon values. After skew synchronization nodes may still operate with different offsets, but because all (one hop) nodes share a frequency standard they are now able to maintain an accurate relative timer for any additional events.

In the second phase we correct for clock offsets. Prior protocols such as NTP and TPSN did this with a two-way message exchange; we take this approach as well, but TSHL considers a *skew-compensated* two-way exchange.



When both phases have been completed we have a model mapping the local, inaccurate clock to the reference time-base. We can then compute a global time for all events, even those before our synchronization, with *post-facto* correction if necessary, similar to RBS.

In phase one, skew is estimated without any knowledge of propagation delay. The quality of our estimate of skew is dependent on the consistency, not the duration, of propagation delay. In our current approach we assume propagation delay is constant over the message exchange. Underwater, changes in temperature and pressure affect the speed of sound and so will change propagation delay over long durations. However, we expect a 20-beacon exchange to take a few seconds; over this period our assumption seems reasonable. Verifying this assumption is an area of future work, however if it does not hold we can fall back on a statistical model of skew as done by Fober et al. [8].

Beacon Nodes can either be specialized nodes with accurate clocks, perhaps connected to an external time reference like a tethered GPS receiver on a buoy. Alternatively, Beacons Nodes could be elected or externally designated.

As noted in Section 2.3 and first addressed TPSN [9], a great deal of non-determinism in message exchange can be removed by placing message timestamping in the MAC layer. For highest accuracy we believe low-level timestamping is essential. We expect our acoustic modems to provide this bit- or byte-level radio access, making MAC-level timestamping easy to implement. Our current prototype implementation includes this feature (as discussed in Section 7.1).

A second assumption of TSHL is that clocks are *short-term stable*. Clock frequency and hence skew, must over a short period of time (typically 5-10 minutes) remain constant. Short term instability occurs mainly due to environmental factors such as sudden variation in temperature, supply voltage or shock [24]. This assumption allows us to model the clock skew using linear regression and use it for predicting the future time accurately as well. One typically accommodates long-term instability by periodic resynchronization (as described in RBS [7]); optimization of that interval is not part of our current research.



**Figure 6. Phases in TSHL**

## 5.2 Details

Figure 6 shows the message exchange in TSHL’s two phases.

In Phase 1 of the protocol, each node in the broadcast range of a Beacon Node models its clock skew. We define the Beacon Nodes clock as the reference timebase (i.e.,  $f_B(t) = \hat{f}_B(t) = t$ ). The Beacon Node then sends out enough Beacon Messages for skew estimation. Our current simulations and implementation send 25, needed for reasonable linear regression. In the future we can make this value adaptive based on feedback from receivers; each receiver can estimate error and enter Phase 2 when skew error reaches a target threshold.

Each Beacon Message  $BM_i$  contains the transmit timestamp  $t_{B,i}$  obtained at the MAC level, just before the message left the Beacon Node. Each receiver  $R$  gets this message at absolute time  $t_{B,i} + D_{B \rightarrow R}$  where  $D_{B \rightarrow R}$  represents the unknown propagation delay between the Beacon Node and the receiver. It then assigns it the local time  $f_R(t_{B,i} + D_{B \rightarrow R})$ . This local time includes error due to clock skew and offset in addition to propagation delay. However, as with Fober et al. [8], we can still model the drift of the local clock with respect to the Beacon’s reference clock by doing a linear regression on the *difference* between receive timestamp and the timestamp in the message. This difference changes by the same amount as clock skew as they drift apart (provided our assumption that path delay is constant over our estimation interval).

Thus, for  $N$  messages  $M_i$ , we do linear regression over the following data points:

$$(t_{B,i} - f_R(t_{B,i} + D_{B \rightarrow R}), f_R(t_{B,i} + D_{B \rightarrow R})) \quad (5)$$

This computation gives us skew correcting conversion of the local time, so from now each node is *skew synchronized* with its skew corrected local time represented as  $f'_R(t)$  (as opposed to a fully synchronized clock  $\hat{f}_R(t)$ ). This skew correction allows us, in phase two, to do correct offset estimation and ultimately map prior or future events to the reference timebase.

Phase 2 is similar to the classical two-way synchronization exchange as shown in TPSN and Figure 4. TSHL’s two-way exchange differs in that we correct for skew when computing the clock offset. When the receiver obtains enough

beacons to estimate the skew it sends a Synchronization Request message with  $T1 = f'_R(T1)$ , the skew-corrected local timestamp. The Beacon Node records its local version of this  $T2 = f_B(T1 + D_{R \rightarrow B})$  and returns this value to the receiver in a Synchronization Reply message timestamped at  $T3$  who computes a skew-corrected receive time  $T4 = f'_R(T3 + B \rightarrow R)$ . Finally the receiver can compute its clock offset:

$$offset_R = [(f_B(f'_R(T1) + D_{R \rightarrow B}) - f'_R(T1)) - (f'_R(T3 + B \rightarrow R) - T3)]/2 \quad (6)$$

When this exchange completes, the nodes are able to factor out the error that occurs because of skew and get the exact offset.

When considering TSHL, it is clear that the first phase uses the RBS/FTSP technique for modeling skew but—since the propagation delay assumption does not hold—not the offset. In the second phase we factor out the offset using the NTP/TPSN 2-way exchange, however, unlike their approach, here we compensate for skew. Thus over all, we cater to the high latency in underwater acoustic networks. In the next section we demonstrate that these changes are required for high synchronization accuracy in these networks.

## 6 Performance Evaluation of TSHL

In this section we present some preliminary results evaluating our time synchronization algorithm. We also describe the goals of our simulations, the methodology, comparison against existing time synchronization protocols for sensor networks, and its dynamics under different parameters.

### 6.1 Goals and Methodology

Our goals in conducting this evaluation of TSHL were two-fold:

- Compare performance of TSHL with other established sensor network time synchronization schemes such as TPSN and RBS.
- Understand the dynamics of TSHL under varying parameters such as node clock skew, granularity, and prediction error.

Although our early simulation results compared TSHL against TPSN, RBS, and FTSP. However, the results in this paper show comparisons only against TPSN. We omit RBS

and FTSP results because those protocols do not consider propagation delay at all and so exhibit error proportional to propagation distance. This error is significant, even at short distances. RBS, for example, shows 6ms error at distances of 10m, and the error grows to 100ms and more at larger distances.

Overall, our simulations show that TSHL maintains low synchronization error, irrespective of the path delays, and its accuracy is directly proportional to the receive jitter (due to bit encoding/decoding or interrupt handling) or the granularity (smallest time increment possible) of the clocks used on the nodes themselves.

## 6.2 Simulation Setup

We simulated the protocols in a custom event driven, packet level simulator designed for an acoustic underwater environment with high latency. Each node's clock was simulated as having some skew and offset relative to the global simulation time. We place a single Beacon Node, with no skew and zero offset in all simulations; all receiving nodes are within a 500m radial distance. We model encoding/decoding and interrupt handling errors referred in [14] by introducing a Gaussian receive jitter. Gaussian distributions have been found to provide reasonable approximations of this error, both by Elson and Estrin [7] above the MAC layer, and more recently by the authors of TPSN [9] even with MAC-layer timestamping.

The simulator allowed us to alter the following parameters in the simulations:

- Number of Beacon Messages broadcast by Beacon Nodes in phase 1 of the protocol.
- Individual clock skew and offset.
- Granularity of the clocks.
- Receive Jitter distribution.

Although we assume the environment was stable during the protocol design, as a worst case we vary water temperature uniformly between 25 and 35°C for each packet sent by the Beacon Node. This temperature variation, with randomized receive jitter, accounts for any inaccuracy observed in the TSHL protocol. In each of the simulation test performed the following parameters were used and are assumed to hold unless specifically mentioned otherwise:

- Skew = 40 ppm.
- Offset = 10 $\mu$ s.
- Number of Beacon Messages = 25.
- Granularity = 1 $\mu$ s.
- Receive Jitter = 15 $\mu$ s.

Each data point shown in a graph is the mean of 1000 simulation runs. Error bars show standard deviations.

## 6.3 Comparative Evaluation

We now compare TSHL performance with TPSN, demonstrating that TPSN accuracy deteriorates at high latencies because it did not need to model skew in its environment. In the first experiment we evaluate error as a function of the distance between the receiver and the Beacon node. To confirm that TPSN error is proportional to clock skew, our second




Figure 7. Comparison of Instantaneous Error




Figure 8. Time Estimation Error

second experiment examines offset error at a fixed distance as clock skew varies.

### 6.3.1 Offset Error Immediately after Synchronization

In this test we measure the *absolute* offset error as a function of distance. (Offset error is the difference between the global time and the *corrected* local time of the node.) First we measure it at the instant immediately preceding the final synchronization exchange (TSHL phase 2).

We expect that the variation in distance (and thus the effect of clock skew) should cause deterioration of error for TPSN, but not for TSHL, as it specifically accounts for that.

Figure 7 confirms that the error increases with receiver distance for TPSN, from about the same accuracy as TSHL at distances less than 100m, to about double the error at 500m. Error does increase for TSHL as well, but by a much smaller



amount; about 12% over 500m. We also observed, in other related experiments (not shown here for brevity), that the synchronization error of TSHL was varying linearly as the mean of the receive jitter distribution.

This simulation presents error immediately after the synchronization ends, thus these results represent the best case performance for TPSN-like protocols that do not model the clock and consider skew. We consider that case next.

### 6.3.2 Time Estimation Error

Here we compare the performance of each of the scheme in predicting the time *after* a particular delay from the time the final synchronization exchange occurred. For this section we consider a receiver placed 400m from the Beacon Node.

Since TPSN does not model the skew and simply measures the offset at that time, we expect it to show a larger error as the time progresses (and the clock skew causes larger drift from the prior synchronized value). TSHL, on the other hand models the clock drift and offset, and hence would be able to correct the local clock more accurately in the future.

Figure 8 confirms our expectations. Error in both protocols is linear with time, but the slope of TSHL is much less because it models skew. Even after 5 seconds, with TSHL, error in offset is below  $50\mu\text{s}$ .

Note that RBS and FTSP do model skew and would do well here. We do not show them because at 400m, their error due to propagation overwhelms the benefits of modeling skew.

(Note that these simulations assume a constant, random clock skew. Over these timescales, variation in clock skew is unlikely.)

### 6.3.3 Effect of Variation of Skew

In this experiment we vary the node skew with respect to the global (and the Beacon Node) clock, and observe the effect on the accuracy of the synchronization. The nodes are kept at a distance of 400m from the beacon node. We vary skew from 5 to 100ppm; this range is wider than typical for real clocks by a factor of about two.

Since TSHL models the skew in phase one, its should be able to cater for whatever the skew is. (Its error should be dominated by mostly by the non-determinism in the delay path i.e. receive jitter, not skew.) TPSN on the other hand, should show increased error as the skew (and thus the error over a single exchange) increases.

Figure 9 validates this expectation as over a wide range of clock skew Synchronization error remains nearly constant for TSHL. Again, this difference is because of TPSN's not modeling skew.

## 6.4 TSHL Parameters

Having established the robustness of TSHL in a high latency acoustic channel, we try to analyze the dynamics of the protocol by altering some of the key parameters of the algorithm. Note that we have not shown results of varying the initial offset, since all protocols are able to factor this out consistently.

### 6.4.1 Number of Beacon Messages

In this simulation we analyze the effect of changing the number of Beacon Messages transmitted in phase 1. Since we do linear regression over these points to estimate clock




Figure 9. Effect of clock skew

skew we expect additional messages to result in higher accuracy.

In this simulation we set the receive jitter to  $30\mu\text{s}$  and consider a receiver 400m from the Beacon node. We show two separate results: One where the error is calculated 0.5s after the final synchronization exchange, and in the other 5s after the exchange.

Figure 10 shows the results of these simulations. First, we observe that there is diminishing benefit in increasing the number of beacons after a certain optimal value. This value seems to be at around 25 beacons. These results corroborate very similar results for the group dispersion as a function of beacons in RBS [7].


In addition, we observe (but not shown here) that at shorter Receiver–Beacon Node distances, fewer beacons are needed to get equivalent accuracy.

### 6.4.2 Variation of Clock Granularity


In this section we observe the effect clock granularity has on the accuracy of the protocol. Since the clock granularity puts a fundamental limit on how accurate a clock can be, we can expect that the performance will degrade with as clock granularity increases. While clock granularity is limited by hardware constraints, software controlled interrupts allow one to select coarser clock granularities to reduce interrupt rate and energy consumption, thus there may be a desire to vary this parameter.

For this test we placed the receiver at a distance of 100m from the Beacon Node node. We examine granularity at uniform increments of  $5\mu\text{s}$ .

Figure 11 shows that both the mean error and the standard deviation of the error increases as the clock granularity degrades. However this is not a bad reflection on the protocol itself; it just goes to show the fundamental limits imposed by the granularity of a node clock on the level of synchrony that can be achieved.



**Figure 10. Effect of Changing the Number of Sync Beacons**



**Figure 11. Varying the Clock Granularity**

## 7 Experimental Evaluation of TSHL

Analysis and simulation help demonstrate the importance of considering propagation delay in acoustic communication, but these approaches necessarily simplify the detail of the real world.

We are in the process of implementing TSHL. This section describes our ongoing implementation. We expect to have preliminary experimental results based on this implementation by the end of May.

### 7.1 Approximating Underwater Acoustics

The main challenge in testing TSHL with high-latency acoustic communication is that our underwater acoustic modems are currently under development. A number of important design issues must be solved before short-range acoustic modems are available. We considered but rejected the alternative of substituting off-the-shelf, long-range

acoustic modems. Not only do they have very different characteristics than our preliminary short-range design, but prior time synchronization protocols have demonstrated the importance of integrating timestamping with the MAC layer [9]. Such integration is impossible with easily available packages.

Instead, we substituted in-the-air acoustic communication for underwater communication. Ultrasound and audible sound have been widely used for localization in surface sensor networks [18, 19, 10]; we adopted the Cricket platform due to its commercial availability and good support for low-level hardware access.

In-the-air acoustics changes several things. First and most importantly, the ratio of the speed of sound in air to water is about 1:5 ( $\sim 300\text{m/s}$  vs.  $\sim 1500\text{m/s}$ ).

This scaling factor plays to our advantage in that shorter distances appropriate for in-office testing scale to much longer equivalent underwater distances. Second, in-the-air acoustics capture some of the noise and random variation of the real world. Although we make no claims that different air temperatures and multipath effects of an office accurately model underwater currents and propagation, it does capture some unpredictable variation.

Our choice of the Cricket platform forced a second approximation. The Cricket hardware consists of a Mica-2-like mote core with Chipcon radio, and it adds a 40kHz ultrasound-based transmitter and microphone [5]. While this platform supports sending and receiving acoustic pulses, it current does not support data transmission. We therefore approximate the acoustic channel by sending an acoustic pulse and a coupled data packet over the standard Chipcon radio. Although the data packet arrives quickly, we delay reception of it until the acoustic pulse arrives, experimentally accounting for propagation latency, and we add a computed transmission delay, accounting for bandwidth limitations.

## 8 Conclusions and Future Work

In this paper we have presented *Time Synchronization for High Latency*, a time synchronization protocol for high point-to-point latency environment. As short-range, underwater sensor networks are developed we expect protocols that consider high-latency communication to be increasingly important. TSHL represents a first such protocol.

Through analysis and simulation we explored TSHL for a wide range of characteristics. Prior protocols like RBS and FTSP exploit the rapid propagation of RF; in an high-latency acoustic network they perform quite poorly (about 6ms error at 10m, growing to hundreds of milliseconds at our target maximum range of 500m). By considering propagation latency, TPSN does much better. However, we demonstrated that it is critical to consider skew as well, even during the synchronization exchange. At short ranges TPSN and TSHL accuracy are equivalent, but at long ranges TSHL shows up to two times better accuracy. An implementation of TSHL in an in-the-air acoustic network is underway.

## 9 Acknowledgments

We would like to acknowledge the help extended by Wei Ye, and Jack Wills in developing the concepts essential

to underwater acoustic communication, particularly in Section 2.1. Also we would like to thank Saurabh Ganeriwal, Jeremy Elson, and Brano Kusy, all authors of previous time synchronization schemes, for extending courteous help in correctly understanding their protocols. The cricket team at MIT, especially David Moore and Michel Goraczko, helped us understand their hardware implementation. Finally Qasim Chaudhari discussed and corrected some of the aspects of the simulations.

## 10 References

- [1] AKYILDIZ, I. F., POMPILI, D., AND MELODIA, T. Underwater acoustic sensor networks: research challenges. *Ad Hoc Networks Journal* (Mar. 2005), to appear.
- [2] ATMEL. <http://www.atmel.com/>.
- [3] BONNET, P., GEHRKE, J. E., AND SESHADRI, P. Querying the physical world. In *IEEE Personal Communications* (2000).
- [4] CHEN, W., HOU, J., AND SHA, L. Dynamic clustering for acoustic target tracking in wireless sensor networks. In *11th IEEE International Conference on Network Protocols (ICNP'03)* (Atlanta, Georgia, USA, November 4-7 2003).
- [5] ELECTRONICS, M. <http://www.mouser.com/>.
- [6] ELSON, J. *Time Synchronization in Wireless Sensor Networks*. PhD thesis, UCLA, 2003.
- [7] ELSON, J., GIROD, L., AND ESTRIN, D. Fine-grained network time synchronization using reference broadcasts. In *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI 2002)* (2002).
- [8] FOBER, D., ORLAREY, Y., AND LETZ, S. Clock skew compensation over a high latency network. In *Proceedings of the ICMC, ICMA San Francisco* (2002).
- [9] GANERIWAL, S., KUMAR, R., AND SRIVASTAVA, M. B. Timing-sync protocol for sensor networks. In *Proceedings of the 1st international conference on Embedded networked sensor systems* (2003), ACM Press, pp. 138–149.
- [10] GIROD, L., BYCHKOVSKIY, V., ELSON, J., AND ESTRIN, D. Locating tiny sensors in time and space: A case study. *Proceedings of the IEEE International Conference on Computer Design* (Sept. 2002). Invited paper.
- [11] HORAUER, M., SCHMID, U., SCHOSSMAIER, K., HÖLLER, R., AND KERÖ, N. Psynutc—evaluation of a high precision time synchronization prototype system for ethernet lans. In *34th Annual Precise Time and Time Interval Meeting (PTTI)* (December 2002).
- [12] KOPETZ, H., AND SCHWABL, W. Global time in distributed real-time systems. 15/89, Technische Universität Wien, 1989.
- [13] LAMPOR, L. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM* 21, 7 (1978), 558–564.
- [14] MAROTI, M., KUSY, B., SIMON, G., AND LEDECZI, A. The flooding time synchronization protocol. In *Proceedings of the 2nd international conference on Embedded networked sensor systems* (2004), ACM Press, pp. 39–49.
- [15] MILLS, D. <http://www.eecis.udel.edu/~mills/ipin.html>.
- [16] MILLS, D. Internet time synchronization: the network time protocol; RFC 1129. *Internet Request for Comments*, 1129 (Oct. 1989).
- [17] NEUMAN, B. C., AND TS'O, T. Kerberos: An authentication service for computer networks. *IEEE Communications Magazine* (Sept. 1994), 33–38.
- [18] PRIYANTHA, N. B., CHAKRABORTY, A., AND BALAKRISHNAN, H. The cricket location-support system. In *Proceedings of the 6th ACM International Conference on Mobile Computing and Networking* (Boston, MA, USA, Aug. 2000), ACM, pp. 32–43.
- [19] SAVVIDES, A., HAN, C.-C., AND SRIVASTAVA, M. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the ACM International Conference on Mobile Computing and Networking* (Rome, Italy, July 2001), ACM, pp. 166–179.
- [20] STOJANOVIC, M. Under water communication. *Wiley Encyclopedia of Electrical and Electronics Engineering* (1998).
- [21] STOJANOVIC, M. *Wiley Encyclopedia of Telecommunications*. John Wiley & Sons, Inc., 2003, ch. Acoustic (underwater) Communications.
- [22] URICK, R. *Principles of Underwater Sound*. McGraw-Hill Book Company, 1991.
- [23] VAN GREUNEN, J., AND RABAEY, J. Lightweight time synchronization for sensor networks. In *Proceedings of the 2nd ACM international conference on Wireless sensor networks and applications* (2003), ACM Press, pp. 11–19.
- [24] VIG, J. R. Introduction to quartz frequency standards. Tech. Rep. SLCET-TR-92-1, Army Research Laboratory, Electronics and Power Sources Directorate, 1992.
- [25] WANG, H., YIP, L., MANIEZZO, D., CHEN, J., HUDSON, R., ELSON, J., AND YAO, K. A wireless time-synchronized cots sensor platform part ii—applications to beamforming, 2002.
- [26] ZHANG, B., SUKHATME, G. S., AND REQUICHA, A. A. G. Adaptive sampling for marine microorganism monitoring. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2004), vol. 2, IEEE, pp. 1115–1122.