Ultra-Low Power Time Synchronization Using Passive Radio Receivers

Yin Chen[†] Qiang Wang[‡] Marcus Chang[†] Andreas Terzis[†]

[†]Computer Science Department
Johns Hopkins University
Baltimore, MD 21218, USA
{yinchen,mchang,terzis}@cs.jhu.edu

[‡]Dept. of Control Science and Engineering
Harbin Institute of Technology
Harbin, P. R. China, 150001
wangqiang@hit.edu.cn

ABSTRACT

Considering its central importance to sensor networks, time synchronization has received extensive attention by the research community. Nevertheless, we argue in this paper that existing approaches introduce undesirable trade-offs. For example, while GPS offers excellent accuracy for outdoor deployments, the high cost and power consumption of GPS receivers make them prohibitive to many applications. Message-passing protocols, such as FTSP, introduce different sets of compromises and constraints. In this paper, we present an inexpensive and ultra-low power (< 100 μA) mote peripheral, we term the *Universal Time Signal Receiver*, that leverages the availability of time signals transmitted by dedicated radio stations around the globe to provide access to UTC time with millisecond-level accuracy. We present experimental results measuring signal availability, quality of synchronization across motes, and power consumption. We show that the proposed universal time signal receiver achieves global time synchronization and for applications where millisecond-level precision is sufficient, it consumes up to 1,000 times less energy than GPS or FTSP.

Categories and Subject Descriptors

C.3 [**Special-Purpose and Application-Based Systems**]: Real-time and embedded systems; B.m [**Hardware**]: Miscellaneous

General Terms

Design, Measurement, Experimentation

Keywords

Time Synchronization, Time Signal, Wireless Sensor Networks, Low-power

1. INTRODUCTION

Time synchronization is a core service for wireless sensor networks used by MAC protocols (e.g., [38, 41, 47]), energy-efficient tree collection protocols [4], distributed computation [8], and applications such as target tracking ([2, 14]) and counter-sniper systems [42]. Considering the importance of this service, multiple mechanisms have been proposed in the literature. Broadly speaking, we can divide these approaches into those that use message passing to synchronize the nodes of a sensor network (e.g., [21, 24]) and those which rely on specialized hardware (e.g., [25, 37, 38]).

We argue that while successful in the context they were proposed, existing mechanisms lack one or more desired properties: low-power consumption, support for large scale networks, accuracy that is independent of network size, access to UTC time, support for disconnected operation in sparse networks, low latency, and ability to operate in both indoors and outdoors environments.

In this paper we present a *universal time signal receiver* that satisfies all these requirements by leveraging the availability of time signals transmitted from dedicated radio stations around the globe. The receiver can be configured to output UTC time streams through a UART interface and a pulse-per-second signal that a mote can use to count the number of elapsed seconds. The universal time signal receiver combines an off-the-shelf radio chip with an ultra-low power microcontroller from Microchip's PIC family that decodes the signal and interfaces with the mote. Due to the simplicity and the low data rate of the time signal, it is possible to select a PIC such that the average power consumption of the receiver as a whole is 70-90 μA. The resultant low-power and low cost ($3.52 in small quantities during the second half of 2010) time signal receiver can easily be part of every deployed sensor mote. Our experimental results indicate that at 2,400 km away from the WWVB station, the signal is available 47% of the time indoors and 75% of the time outdoors, while at 700 km away from the DCF77 station, the signal is available 97% of the time indoors. Furthermore, we find that the proposed receiver achieves accuracy in the order of 2 msec. Although GPS and message passing based synchronization protocols (with short synchronization interval) can achieve higher synchronization accuracy, they also incur higher power consumptions. For applications that require millisecond or lower timing accuracy, our universal time signal receiver is well suited and consumes up to 1,000 times less energy.

This paper makes the following four contributions: *(1)* we describe how the time signals can be used to provide global synchronization in a sensor network and present a low-power and low-cost *Universal Time Signal Receiver* that interfaces to existing motes. *(2)* We present extensive results about the availability of the WWVB and DCF77 signals in multiple environments and measure the power consumption as well as the accuracy that our receiver achieves. *(3)* We compare our time signal receiver with message passing based synchronization protocols and GPS receivers, and identify the advantage of using the time signal re-

IPSN'11, April 12–14, 2011, Chicago, Illinois.

Station	Country	Frequency	Launched
MSF	Britain	60 kHz	1966
CHU	Canada	3330, 7850, 14670 kHz	1938
BPC	China	68.5 kHz	2007
BPM	China	5, 10, 15 MHz	1981
TDF	France	162 kHz	1986
DCF77	Germany	77.5 kHz	1959
JJY	Japan	40, 60 kHz	1999
RBU	Russia	66.66 kHz	1965-74
HBG	Switzerland	75 kHz	1966
WWV	USA	2.5, 5, 10, 15, 20 MHz	1920's
WWVB	USA	60 kHz	1963

Table 1: Some of the time signal radio stations.

ceiver when targeting millisecond or lower synchronization accuracies, as a result of the combination of fluctuating clock skew and extended synchronization interval. *(4)* We describe how sensor network protocols can leverage the availability of inexpensive global synchrony to improve energy efficiency. Specifically, we present a drop-in replacement for the existing TinyOS LPL MAC and show how by using our replacement the duty cycle of a CTP application decreases by 83%. We also show how the proposed universal time signal receiver could be used to maintain the time state for energy-scavenging motes.

This paper has seven sections. The section that follows reviews signals used to disseminate time information and elaborates on the time signals transmitted by the WWVB and DCF77 radio stations. Section 3 frames our proposal in the context of the extensive work done in time synchronization for wireless sensor networks. In Section 4 we elaborate on the design of the low-power time signal receiver that we developed and in Section 5 we present results regarding the availability of WWVB and DCF77 signals in different environments, the power consumption of our time signal receiver, and the accuracy that it achieves. Section 6 outlines how one can leverage the availability of global synchrony to reduce WSN energy consumption and we conclude in Section 7.

2. PERVASIVE TIME SIGNALS

Various signals have been created throughout history to serve the purpose of disseminating time information, evolving from sound and visible light to radio transmissions with increasingly sophisticated modulation techniques. Today, GPS is by far the most well known radio frequency signal that, while mainly used to perform localization, provides very accurate time information. Moreover, dozens of radio stations across the world are dedicated to broadcasting time information. Table 1 presents a non-exhaustive list of such time signal stations and their operating frequencies. Most of the stations transmit in the low (LF) to high (HF) frequency radio bands using amplitude modulation (AM). This modulation scheme simplifies the reception and decoding of the time signal, thus enabling inexpensive and energy-efficient receivers. Taking the distance to the radio station into account, the time-of-flight delay can be subtracted from the decoded time signal, in order to achieve global time synchronization. In this paper we focus on two LF time signals, WWVB and DCF77, which cover most of North America and Europe, respectively.

Propagating LF radio waves consist of two components: a ground wave and a sky wave. The ground wave interacts with the Earth's surface and thus follows the Earth's curvature while the sky wave refracts on the Earth's ionosphere enabling it to travel greater distances. Specifically, for the 50 kW DCF77 transmitter this results in three distinct regions: within 600 km from the transmitter, the ground wave dominates signal reception, while beyond 1100 km only the refracted sky wave can be received. In the 600-1100 km region, the ground wave and sky wave are of equal magnitude and both positive and negative interference can occur. The refractive capabilities of the ionosphere are mainly determined by solar activity, and thus follows diurnal, seasonal, and solar cycles.

The WWVB radio station is located near Fort Collins, Colorado, operated by the National Institute of Standards and Technology (NIST) [22]. The WWVB time signal modulates a one-bit-per-second time code onto a continuous 60 kHz carrier wave via pulse width modulation with amplitude-shift keying, as explained below.

Each WWVB data frame contains 60 bits and therefore lasts exactly one minute. Data frames are transmitted back-to-back at a rate of 60 frames per hour or 1440 frames per day. Each data frame starts exactly at the top of a UTC minute and each of the 60 bits starts at the exact top of each corresponding UTC second. At the beginning of each second, the power of the 60 kHz carrier wave is reduced by 17 dB for a duration of 0.2, 0.5 or 0.8 seconds, corresponding to a 0, 1 or "marker" bit respectively. Figure 1 illustrates the content and format of one WWVB data frame, which encodes the time information of 06:11 UTC on the 144^{th} day (May 24) of 2010. One can see from the figure that time is represented as binary-coded decimals. Accordingly, a time signal receiver can detect the fall in the carrier wave energy as the start of a UTC second and use its local timer to record time stamps of those events. After receiving a whole data frame, the receiver extracts the UTC time from the decoded bit stream. Combined with the start-of-a-second time stamps, the receiver can calibrate its local clock to UTC time.

The German national meteorology institute (PTB) transmits the DCF77 time signal across Europe [35]. The signal is transmitted from Mainflingen (25 km from Frankfurt), Germany, using a 77.5 kHz carrier frequency. Although the DCF77 signal also contains the current day, time, and year, similar to WWVB, there are some subtle differences. For example, the data frame only contains 59 bits, but unlike WWVB, all 59 bits are used. Specifically, the fields encoding the hours, minutes, and date are each protected by a parity bit in order to detect data corruption and 14 of the bits can even be used for emergency warnings and weather forecasts. The time itself is reported in Central European Time (CET) and the time signal is encoded using both analog and digital modulations. Analog amplitude modulation (AM) is performed by reducing the signal strength by 6 dB for 100/200 ms to encode a 0/1 value. After the 59th bit the signal strength stays constant for a whole second to signify the beginning of a new minute. For the digital encoding, the time code is first multiplied by a pseudo-random sequence and then encoded onto the carrier signal using a $\pm 13^{\circ}$ phase-shift keying (PSK) [18]. The end result is a time signal in which the 59-bit data frame spans 793 ms, meaning an entire data frame can be transmitted every second. Combined with higher phase modulation accuracy, the digital signal provides both faster time synchronization and higher precision than the analog signal.

NIST states that the 60 kHz carrier wave for WWVB has a frequency uncertainty of 1 part in 10^{12} and WWVB can provide accurate UTC time with 100 μs uncertainty [22, 32]. With its 70 kW transmitter WWVB covers most parts of the United States with a signal field strength $\geq 100\mu V/m$. Moreover, NIST is considering adding another WWVB radio station at the east coast of the United States to further increase coverage [36].

The PTB [35] offers similar frequency precision, but claims that the DCF77 can provide accurate CET with 5.5 μs uncertainty. The 50 kW transmitter covers Central Europe completely, and even

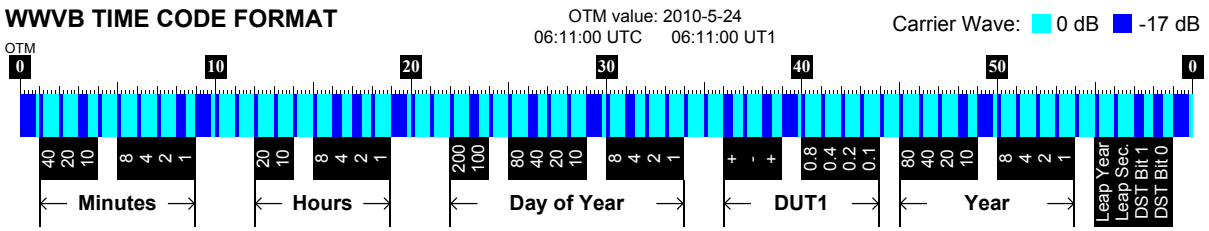

Figure 1: WWVB time code format. A 60-bit frame is transmitted at the top of a UTC minute and each of the 60 bits starts at the beginning of each second. Individual bits are encoded using pulse width modulation. This example frame encodes 06:11 UTC on the 144^{th} day (May 24) of 2010, and the DST bits indicate that the daylight saving time is active.

reaches Moscow and Northern Africa with a signal field strength $\geq 100\mu V/m$. Only the northern parts of Scandinavia, above the Arctic Circle, are not covered.

3. RELATED WORK

GPS is an obvious choice when considering time synchronization protocols based on specialized hardware. Especially for sensor networks deployed outdoors (e.g., [16, 20]) the reception conditions are favorable. Due to the strict timing requirements in measuring the time-of-flight of radio signals, the GPS can be used as reference to UTC global time with nanosecond precision and immunity to partitioning of the sensor network. However, since the current draw of GPS receivers is still in the 20-40 mA range ([25]), and it takes between 30 seconds and up to several minutes to update the satellites' ephemeris every four hours, using the GPS for real-time clock synchronization can significantly impact a network's energy budget. Furthermore, because the high-frequency GPS signal is transmitted over vast distances, it is more susceptible to attenuation caused by solid objects, severely limiting the use of GPS indoors. On the other hand, applications that perform post-mortem reconstruction of collected timestamps, can apply the ephemeris during post-processing and can therefore achieve synchronization by turning on the GPS receiver for only 200 ms at a time [25].

Huang et al. reported that Zebranet nodes experienced loss of GPS signal when the zebras hid beneath treetops [20]. On the other hand, time signals from RF stations do not require line-of-sight reception due to the higher penetration of long-wave radio signals and their ability to travel along the Earth's curvature. Gupchup et al. used a single GPS-equipped mote as an anchor to estimate the global time at all the other motes in their sensor network [16]. Because of the high power consumption, this anchor mote was connected to the electricity grid during the experiment. Considering the low monetary cost and power consumption of our universal time signal receiver, it is conceivable that every node in a sensor network could be equipped with such a receiver.

Because of the higher penetration properties of sub-gigahertz frequencies (compared to those used in GPS, WiFi, and ZigBee) and the low power consumption of amplitude modulated signals, low- and medium frequency AM signals have also been popular in sensor networks, both for time synchronization [1, 34, 37, 38] and out-of-band signaling [7, 9, 15].

Rowe et al. used a commercial AM transmitter to transmit synchronization pulses to the FireFly network nodes deployed inside the building [23, 38]. These schemes used the building's power lines as an antenna for the AM transmitter. Each FireFly mote is equipped with an AM receiver to synchronize itself with the transmitter. In turn, this synchronization is used for an efficient TDMA protocol. The AM receiver achieves sub-20 μs synchronization uncertainty but draws 5 mA of current. Moreover, the AM receiver has to be powered on all the time in order to keep counting the synchronization pulses. The universal time signal receiver described in this work uses dedicated hardware to sample and decode the time signal and consumes two orders of magnitude less energy. Furthermore, it can extract UTC time from a single time signal data frame and can therefore be safely duty-cycled without losing track of global time. We note that although Rowe et al. [23, 38] proposed using WWVB receivers for outdoor networks, they did not report any performance data. Also, the WWVB receiver used in that study consumes an order of magnitude more energy than our universal time signal receiver.

Rowe et al. [37] also developed a hybrid scheme, based on both specialized hardware and message passing. Building upon their previous work, they developed a radio receiver specifically tuned to the 60 Hz alternating current from common household power lines. By using a phase-locked-loop (PLL) to lock on to the frequency and averaging the oscillations over time, their time board is able to generate a pulse-per-second signal whose uncertainty is less than a millisecond. With this stable clock source they use a variant of TPSN to establish a phase offset between neighbors. While ingenious, this scheme is limited to deployments within buildings.

Halpern et al. proposed the first message passing synchronization protocol that used the message broadcast principle to synchronize computers connected through Ethernet [17]. This technique was later extended to sensor networks by Elson et al. [12]. The proposed RBS protocol establishes local time synchronization by using a reference node to broadcast message beacons. Participating nodes subsequently exchange records of the reception times and calculate an offset that minimizes the error of the local clock. However, the overhead of exchanging these timetables increases quadratically with the number of participating nodes. Dai et al. addressed this lack of scalability by having the reference node perform a single round-trip-time measurement to determine its own transmission delay and subsequently broadcast this delay to all the receivers [10]. This technique achieved 29.5 μs precision in a 3-hop network. Instead of measuring the RTT, Maróti et al. use MAC layer time stamping on both the sender and receiver nodes in FTSP [24]. Combining MAC-layer timestamping with clock skew compensation through linear regression, FTSP achieved 3 μs precision in a 6-hop network.

Compared to hardware-based mechanisms, message passing protocols require network nodes to be connected. Specifically, a node must be both within communication range of the transmitter and have its radio on to receive the time update. Furthermore, it was found that FTSP errors increase exponentially with network diameter, leading to significant errors even in mid-size networks with 10-20 nodes [21]. Lenzen et al. developed PulseSync to mitigate this effect by rapidly disseminating time information [21]. Fi-

message	'PGT','yyyymmdd','hhmmss','s','cs*'
'PGT'	header
'yyyymmdd'	year, month, day
'hhmmss'	hour, minute, second
's'	status byte indicating the signal quality and internal state
'cs*'	check sum

Table 2: The date and time message definition.

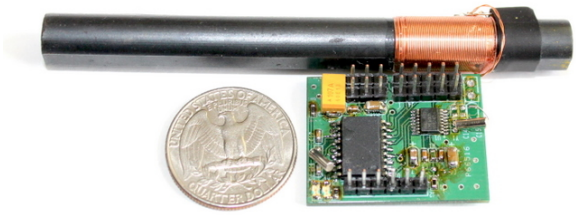

Figure 2: Prototype of the proposed universal time signal receiver. The ferrite antenna is tuned to 60 kHz to enable the reception of the WWVB signal.

nally, Schmid et al. analyzed and quantified the adverse effects of clock skew variations and long message exchange intervals on the accuracy of message passing protocols [40]. They found that synchronization uncertainties can deteriorate to millisecond range when operating at synchronization intervals longer than 500 seconds. Schmid et al. also proposed decoupling the clock distribution tree from the routing tree to prevent the motes with more dynamic clock skew fluctuations (e.g., due to environmental factors such as temperature) from propagating synchronization errors downwards [39].

4. DESIGNING A LOW-POWER UNIVERSAL TIME SIGNAL RECEIVER

This section presents the design of our universal time signal receiver. We analyze the requirements of such a hardware device and discuss our design choices and implementation details.

4.1 System Requirements

Low power consumption, high accuracy, low cost, and small form factor. These are all fundamental requirements when designing motes and by extension important design requirements for our time signal receiver as well. Specifically, low power consumption will enable liberal use of the receiver even in sensor networks with tight energy budgets, while striving for high timing accuracy is an inherent goal for all timing modules. There is however a trade-off between power consumption and timing accuracy. We decided to settle for millisecond level accuracy by leveraging the pervasive time signals introduced in Section 2, because these AM-modulated time signals can be received with inexpensive and low-power radios. Furthermore, we observe that millisecond level accuracy is actually sufficient for many sensor network applications, especially in the area of environmental monitoring, since these deployments typically measure slow changing physical values [16]. Finally, low cost and small form factor would increase the feasibility of widely deploying this universal time signal receiver.

4.2 System design

In order to achieve low power consumption and small footprint, we based our universal time signal receiver on the CME6005 radio chip from C-MAX [6]. We chose the CME6005 for two reasons. First, its frequency range spans from 40 to 120 kHz, which enables us to receive the WWVB, DCF77, JJY, MSF and HBG time signals. Second, it offers low power consumption: less than 90 μA in active mode and 0.03 μA in standby mode. Furthermore, with a sensitivity of 0.4 μV the CME6005 also has excellent reception capabilities. The CME6005 radio chip can receive, demodulate, and convert analog time signals into digital outputs on a CMOS/TTL I/O pin. For amplitude modulated time signals, this means that when the amplitude is high/low, the output value of the I/O pin is also high/low.

Although this signal could be connected directly to a mote's GPIO pin and be decoded by the on-board MCU, doing so would add additional complexity and overhead on the software running on the mote. Furthermore, by relying on the mote's MCU to sample and decode the signal, the accuracy becomes dependent on the system's current load, while the MCU's active time and therefore power consumption increases. Hence, we decided to use a separate MCU to process and decode the time signal and serve as a dedicated standalone timing module.

We choose the extreme low-power microcontroller PIC16LF1827 from Microchip [29] for decoding time frames. This PIC consumes 600 nA in sleep mode with a 32 KHz timer active, and 800 μA when running at 4 MHz. By connecting the CME6005's output pin to an interrupt-enabled input pin on the PIC, we allow the reception and decoding of time frames to be interrupt-based. Considering the 32 KHz timer, the PIC can measure the interval between interrupts and at the same time maintain a local clock with a resolution of approximately 30 μs. In order to achieve low-power operation, we set the PIC to sleep mode with the 32 KHz timer running. When an edge from the CME6005 arrives, the corresponding interrupt service routine is triggered. For a falling edge, the service routine simply records the 32 kHz timer value and when the rising edge arrives, the PIC compares the current local time to determine the value of the received bit. When a full time code frame has been received, the UTC time can be decoded by following the format definitions, and the decoded time will be stored in memory as the global time.

In addition to the external interrupts triggered by the falling and rising edges, the PIC's timer can also be configured to generate an internal interrupt. We utilize this function to generate a standard one-pulse-per-second (1PPS) signal. By changing the values in the compare register, we can calibrate and compensate for the frequency skews of the 32 KHz crystal by comparing with the decoded UTC time.

Table 2 shows the UART output of the universal time signal receiver, which mimics the NMEA (National Marine Electronics Association 0183) standard for date and time messages. We choose this format for improved compatibility and easier replacement of GPS modules. In addition to the date and time information, the receiver also outputs the signal quality information.

Figure 2 shows a prototype of this hardware design. In this picture, a ferrite antenna tuned to 60 kHz is connected to the CME6005 enabling the reception of the WWVB signal. The overall footprint of the receiver board is 25 mm by 33 mm, with the antenna being its largest external component. The antenna shown in Figure 2 is 100 mm in length. Smaller ferrite antennas (e.g., 60 mm and 23.5 mm in length) are available, but we have not tested them for our prototype.

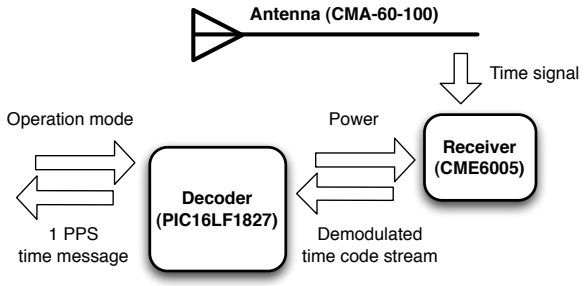

Figure 3: Schematic of the time signal receiver.

Report Interval	Average Current Draw
1 sec	92.6 μA
1 min	69.7 μA

Table 3: Average current draw for the universal time signal receiver ($V_{cc} = 2.88\ V$).

Furthermore, the availability of radio controlled watches (e.g., Casio Waveceptor series) indicates that using smaller antennas is indeed possible. Figure 3 shows the schematic of the prototype in Figure 2. The total cost of the receiver was $ 3.52 in small volumes during the second half of 2010 (CME6005: $1.372, PIC16LF1827: $1.26, antenna: $0.893), comparing favorably to the cost of other hardware-based time receivers.

5. EVALUATION

Next, we evaluate the power consumption of the universal time signal receiver, compare it to synchronization protocols based on message passing and GPS receivers, and quantify the quality and availability of both the WWVB and DCF77 time signals, based on field experiments conducted at multiple geographical locations.

5.1 Universal Time Signal Receiver

The two major power consumers on the receiver board are the CME time signal receiver chip and the PIC microcontroller. As stated in the previous section, the CME chip draws a constant current (less than 90 μA), while the current draw of the PIC in sleep and active mode is 600 nA and 800 μA, respectively. However, when we measure the power consumption of the entire receiver the on-board capacitors (connected between the V_{CC} pin on the PIC and the ground) have a dampening effect on the contribution from each individual component. Furthermore, depending on the residual charge left in the capacitors multiple instances of the same operations can have different power consumptions.

Figures 4(a) and 4(b) show the power consumption trace of the receiver board while outputting the NMEA timestamp every minute and second, respectively. In Fig.4(a) one can see three distinct contributions to the power consumption: first the base level, corresponding to the active CME and sleeping PIC. Second small spikes every second, which correspond to the PIC waking up and running interrupt service routines. Note that we only see a small spike due to the capacitors discussed above. Finally the large spike every minute, corresponds to the PIC decoding the time signal and sending the timestamp over the UART. Because the duration of this operation is long enough to deplete the capacitors we see an actual full increase in power consumption. In Fig.4(b) the timestamp is transmitted every second instead of every minute, which explains why the previously small spikes have turned into large ones.

As the traces show, due to the low bitrate of the time signal, and the simple processing needed, the main contribution to the PIC's power consumption comes from sleep mode. Table 3 summarizes the average power consumptions of the receiver board for the two UART output frequencies. In both cases, the average power consumption is below 100 μA. However, because the 1-pulse-per-second signal on the receiver board is directly synchronized with the one-second-pulse from the time code signal, regardless on whether or not a full time frame has been received, the receiver board can be aggressively duty-cycled: after a time signal has been successfully decoded, the CME receiver can be powered down with the PIC relying on its internal clock (similar to regular message passing time synchronization protocols). The major difference comes when the estimated clock drift has grown too large. Instead of initiating resynchronization through message passing, the PIC only needs to turn on the CME receiver long enough to receive a couple of 1-pulse-per-second pulses in order to correct its clock drift.

5.2 Comparison with Message Passing Protocols and GPS Receivers

As discussed in Sections 1 and 3, the sensor networks literature contains many message passing protocols for time synchronization. For example, protocols such as the widely-used FTSP [24] and the more recent and state-of-the-art PulseSync protocol [21] can synchronize multi-hop sensor networks to microsecond level accuracy. However, the synchronization accuracy of these protocols critically depends on the message exchange frequency and fluctuations in clock skew. For example, both FTSP and PulseSync set their message exchange interval T to 30 seconds. With such high message exchange rate, we expect their power consumption to be significantly higher than that of the universal time signal receiver.

In fact, a first-order analysis of the power consumption for the FTSP protocol shows that maintaining microsecond precision has its toll on the power budget even when nodes use low-power listening (LPL) [3]. According to the FTSP specification, after a node first boots it needs to receive eight synchronization messages from another node in order to be synchronized. Afterwards, the node will transmit one broadcast message during every synchronization period (message exchange interval). This requires the node to keep its radio on for τ seconds in order to send the broadcast, where τ is the LPL sleep interval. At the same time, the node must also listen on its own radio every τ seconds in order to check for channel activity. This check takes approximately 10 ms in the current LPL implementation in TinyOS 2.x.

Disregarding the actual reception time of these messages and considering only the energy that each node spends on its own packet transmissions and checking for channel activity, the average current draw of the radio (assuming it consumes 20 mA in transmit mode [45]) is:

$$I = 20\text{mA} \times (\tau/T + 0.010\text{s}/\tau) = 20(\tau/T + 0.010\text{s}/\tau)\text{mA},$$

where T is the message exchange interval and the radio checks channel activity for 0.010 seconds. The minimum value of I is given by

$$\begin{aligned} I &= 20(\tau/T + 0.010\text{s}/\tau)\text{mA} \\ &\geq 20 \times 2\sqrt{\tau/T \times 0.010\text{s}/\tau}\text{mA} = 4 \times \sqrt{1\text{s}/T}\text{mA} \end{aligned} \quad (1)$$

where the optimality is reached when $\tau = \sqrt{0.010\text{s} \times T}$.

It is obvious from the equations above that choosing a larger

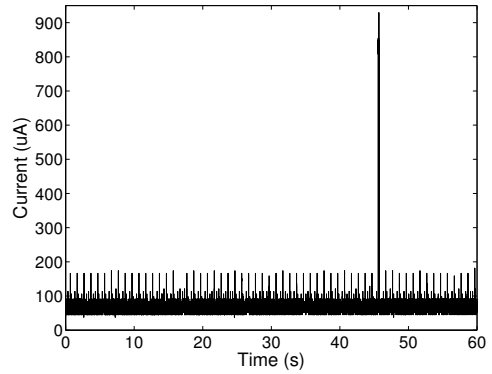

(a) Receiver set to output UTC timestamp every minute.

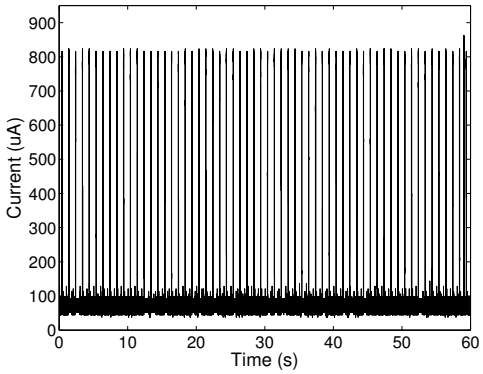

(b) Receiver set to output UTC timestamp every second.

Figure 4: Power consumption traces of the universal time signal receiver while outputting the NMEA timestamps.

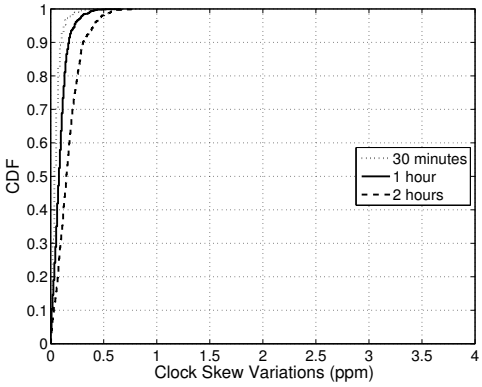

(a) Indoor environment.

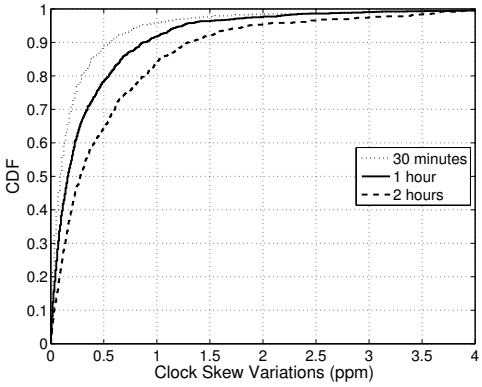

(b) Outdoor environment.

Figure 5: Measured deviations of relative clock skews across a network of nine Tmote Sky motes as a function of time elapsed between synchronization events.

message exchange interval T immediately reduces the power consumption. However, since these messages are used to estimate the relative clock skew to a reference node and these clock skews vary with environmental conditions such as temperature, increasing the message exchange interval also increases the synchronization error due to clock skew variations. Consequently, the larger the message exchange interval T, the bigger the error. For example, a 1 ppm change in clock skew can lead to 1 μs of drift in one second, 60 μs of drift in one minute, and 3.6 ms of drift in an hour, respectively.

The remaining question then is: *how dynamic is the clock skew?* In order to answer this question, we performed a multi-day experiment with motes placed both indoors and outdoors. Figures 5(a) and 5(b) show the cumulative distributions of the differences in relative clock skews as a function of T. For each synchronization interval T, we calculate the average relative clock skew of each node with respect to a reference node. Next, we compute the differences in clock skews between two consecutive synchronization intervals. We used three values for T: 30 minutes, 1 hour and 2 hours. We clearly see that motes in the outdoor experiment experienced significantly higher variations in clock skew when compared to the indoor experiment. This is not surprising since outdoor environments have more dynamic variations in both temperature and humidity, both of which cause fluctuations in clock frequencies. Consequently, message passing protocols must rely on more frequent message exchanges in order to achieve the same level of accuracy for outdoor deployments as for indoors.

Quantitatively, we find that the average difference between two consecutive hours is approximately 0.09 ppm and 0.36 ppm for the indoor and outdoor experiment, respectively, while the maximum difference is 0.67 ppm and 6.68 ppm, respectively. We note that in both experiments, the motes were placed in close proximity because of the wires interconnecting them. In the outdoor experiment, the motes were placed under trees that shaded the direct sunshine and dampened the temperature fluctuations. Therefore, motes that are located farther away from each other, as is the case in environmental monitoring sensor networks, are expected to have more severe clock skew fluctuations than those shown in Figure 5(b).

In order to compare the power consumption of message passing protocols to the universal time signal receiver we perform a similar current consumption analysis. After it boots, the mote needs to receive an entire frame before it becomes synchronized. Subsequently, instead of receiving an entire data frame, the receiver only needs a start-of-a-second pulse to calibrate for its local clock skew. Being conservative, we assume that the PIC needs to turn on the CME for 5 seconds to capture a pulse. Also, we use the results from Table 3 and assume the receiver consumes approximately 70 μA in active mode. Hence, the average current draw for the univer-

Interval (s)	I (mA)	I_{ts} (mA)	Ratio
1000	0.126	3.5e-4	361:1
3,600	0.067	9.72e-5	686:1
7,200	0.047	4.86e-5	970:1
10,800	0.038	3.24e-5	1188:1

Table 4: Comparison of average current draw for FTSP (I) and the time signal receiver (I_{ts}).

sal time signal receiver can be written as:

$$I_{ts} = 0.070\text{mA} \times 5\text{s}/T = 0.35\text{mA} \times 1\text{s}/T \quad (2)$$

We note that (2) does not consider the power that the PIC consumes while in sleep mode (with its 32 kHz timer running). The reason for doing so is as follows. Here we are comparing the power consumption of the universal time signal receiver against running FTSP on the mote. Note that the mote that is running FTSP also draws power to keep its timer running, and this power consumption (54.5 μA in the case of Tmote Sky[30]) is most likely going to be higher than the PIC's 600 nA. On the other hand, if we connect the universal time signal receiver to a mote, the PIC can be powered down by the host mote and be turned on only when the host mote wants to receive time signals. Therefore in Equation (1) and (2) we only consider the power consumptions directly associated with radio activities.

Table 4 compares the power consumption of FTSP to that of the universal time signal receiver when operating at millisecond level synchronization accuracy. In this regime, the time signal receiver consumes two to three orders of magnitude less power than FTSP.

Note that when assuming a change of 1 ppm in clock skew, the uncertainties accumulated over the intervals listed in Table 4 ranges from 1 ms to 10.8 ms. As an example, the experimental results presented in [40] show that seven closely located TelosB motes experience 1 ms synchronization error when using a 500 s update interval, and the error goes up to approximately 7 ms when using 1000 s as the update interval.

In summary, fluctuating clock skew is the fundamental challenge in achieving accurate time synchronization. Message passing based synchronization protocols periodically exchange packets to establish reference points between local time and global time. The accuracy of these reference points depends upon the network topology, the underlying radio hardware, the local clock resolution, and the message passing protocol. Regardless of the quality of these reference points, however, local clock drifts and therefore synchronization error accumulates between reference points. As the synchronization interval (i.e., the time between two reference points) increases, the synchronization error also grows accordingly, which eventually becomes the performance bottleneck of any synchronization protocol.

The issue associated with the combination of clock skew and synchronization interval also applies to the case where each mote employs a GPS receiver to acquire the global time reference. Although GPS receivers can provide global time reference with nanosecond accuracy, if the mote queries the GPS receiver once per second, its local clock can drift by up to 1 microsecond (assuming 1 ppm fluctuation) in between the queries. Note that 1 ppm fluctuation is quite unlikely to occur within one-second time window, nevertheless, the deterioration continues as the interval increases, and can soon reach into the millisecond region. This corresponds to the scenario in which motes are duty-cycling the GPS receivers to save energy as the application requires looser timing accuracy.

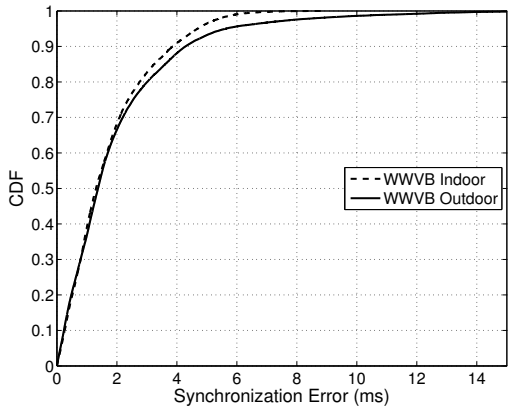

Figure 6: WWVB synchronization accuracy for both an indoor and outdoor experiment.

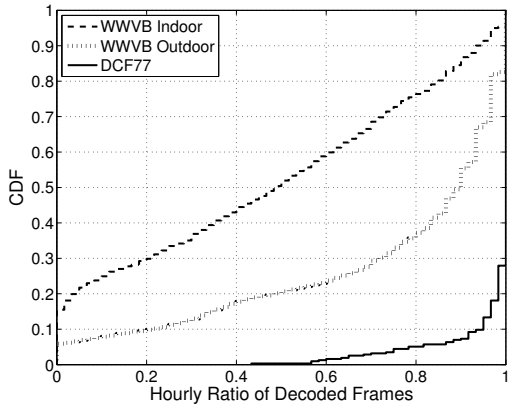

Figure 7: Availability of the WWVB and DCF77 time signal.

On the other hand, the universal time signal receiver provides global time references with only millisecond level accuracy (Section 5.3). While this limitation prevents sub-millisecond accuracy, as the synchronization interval increases, the overall synchronization error becomes increasingly dominated by the clock drifts that accumulate during each interval. Eventually, when aiming for millisecond- or second-level synchronization, larger intervals can be scheduled and synchronization error is completely dominated by local clock drifts during synchronization intervals. Consequently, replacing the universal time signal receiver with a GPS receiver will only have a marginal effect on synchronization accuracy in these cases, while dramatically increasing power expenditure.

5.3 Quality and Availability of Time Signals

In order to quantify the quality and availability of time code signals we performed multiple experiments spanning several days, at different distances to the radio signal source in both indoors and outdoors settings.

We tested the WWVB signal availability at sites located more than 2,400 km away from the signal source. Although by design the WWVB radio station should be capable of covering locations that are even farther away, our testing sites are already challenging the limits of the WWVB signal transmitter, mainly because only the sky wave can be received making reception vulnerable to changing solar activity. On the other hand, the DCF77 signal was tested only 700 km from the signal source. At this distance, mainly the ground

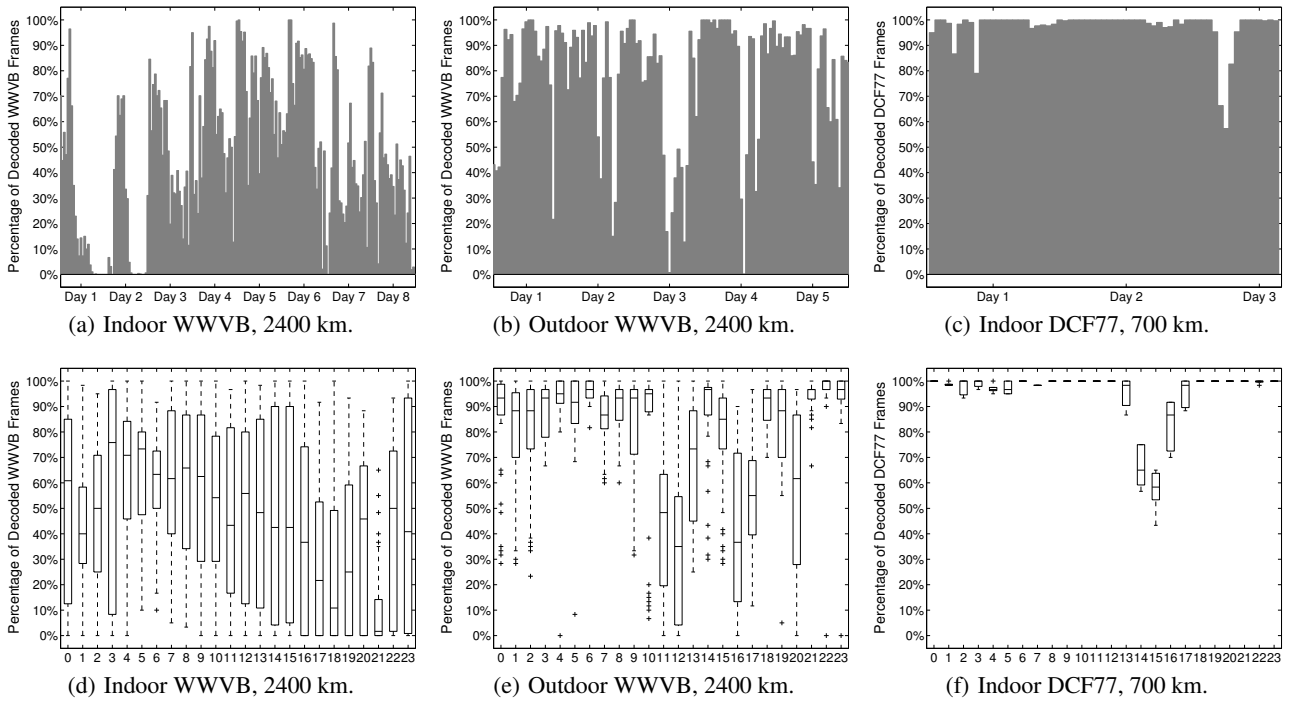

(a) Indoor WWVB, 2400 km. (b) Outdoor WWVB, 2400 km. (c) Indoor DCF77, 700 km.

(d) Indoor WWVB, 2400 km. (e) Outdoor WWVB, 2400 km. (f) Indoor DCF77, 700 km.

Figure 8: Time signal availability across different environments.

wave is received with intermittent interference from the sky wave. In the case of WWVB, we did both an eight-day indoor experiment with eight receivers inside a big apartment building, and an five-day outdoor experiment with nine receivers in a small park near a residential neighborhood. For the DCF77, we did a two-day long indoor experiment with five receivers in an old office brick building. Generally, we observe that besides weather and distance, other sources can interfere with signal reception as well. Specifically our experiments show that: (1) steel frame buildings completely shield the time code signal, (2) brick buildings allow signal reception, (3) CRT screens can interfere from 5-10 meters distance, and (4) laptops can interfere within one meter.

In order to quantify the quality of the WWVB time code signal we connected all the motes to a common signal source with wires. When all universal time signal receivers were synchronized, a pulse was transmitted over the wire every sixth second, signaling the motes to record their perceived UTC time derived from the WWVB signal. For each event, we calculate each mote's UTC difference to the mean for that particular instance and plot the cumulative distribution in Figure 6. Not surprisingly, we again find that the indoor experiment is more precise than the outdoor. However, when compared to the significant differences observed between indoor and outdoor clock skews as discussed above and seen in Figures 5(a) and 5(b), the difference in precision between an indoor and outdoor WWVB signal is relative small. This means that the universal time signal receiver is more robust against clock skews in outdoor deployments and that hybrid deployments in which motes are placed both indoors and outdoors will experience the same level of accuracy. Specifically we find that the 50%, 80% and 90% synchronization error percentiles for the indoor experiment, were respectively: 1.3 ms, 2.8 ms, and 3.9 ms. For the outdoor experiment the error for the same percentiles were: 1.4 ms, 3.0 ms, and 4.3 ms.

Next we explore the availability of the WWVB and DCF77 signals both at different distances from the radio signal and indoor/outdoor locations. We use the ratio of correctly decoded time code frames during each hour. Since each data frame is one-minute long, exactly 60 frames are transmitted every hour. We continuously collected time code frames for several days at the three locations mentioned above and show the cumulative distribution of the ratios of correctly decoded time frames in Figure 7. First, the difference in availability between the indoor WWVB and outdoor WWVB experiments is significant. Specifically, for the indoor experiment we see that during 52%, 85%, and 93% of the hours, the motes were able to decode less than 50%, 80%, and 90% of the available time code frames. For the outdoor experiment, these numbers decrease significantly to 20%, 37%, and 55%. Second, looking at the DCF77 ratios, although also collected indoors, being only a third of the distance away from the radio source, the increase in availability is significant. Specifically, during 1%, 6%, and 7% of the hours the motes were able to decode less than 50%, 80%, and 90% of the available time code frames, respectively.

Finally, we evaluate the availability of the WWVB and DCF77 signals during the course of the experiment. Rather than looking at the cumulative distribution, we consider the hourly ratio of decoded time frames. Figure 8 (a)-(c) illustrates the hourly decode ratio averaged across all receivers. In the case of the indoor WWVB experiment (Fig.8(a)), we see that the signal quality fluctuates significantly, and during several hours, the receivers were not able to decode any time code frames. For the outdoor WWVB experiment (Fig.8(b)), the signal availability overall improves drastically but complete blackouts still occur, albeit less frequently. Last, the DCF77 experiment (Fig.8(c)) shows close to full availability during the course of the experiment.

Figures 8 (d)-(f) present the box plots of the hourly decode ratio for all the motes. Considering the sources of interference, for the WWVB experiments, Figures 8(d) and (e) show that the deviation within each hour is high, meaning the nodes lost reception independent of each other and therefore probably due to local conditions.

On the other hand, when considering the DCF77 experiment in Figure 8(f), the decrease of decoded frames between hours 13 and 17 happens simultaneously across all motes and with the same intensity, meaning the interference is caused by some global phenomena, possibly interference from the sky wave.

5.4 Use with Local Time Signal Generators

The results presented thus far show that each time signal can cover a very large area, possibly an entire continent. Nevertheless, signal strength decays with distance, causing the availability to decrease. As a consequence, there are places with limited or even no time signal availability. As we also saw, it is more difficult to receive the radio signals inside buildings than outside. To accommodate locations with bad or no reception, a local time signal generator that can replicate and/or amplify the original time signal can be deployed. Commercially available, off-the-shelf time signal generators, such as the TSG100 [5], can generate DCF77, JJY60, MSF and WWVB time signals covering areas with up to 50 meters radius. This signal generator complies with FCC regulations and does not require a license to operate. We performed a 10-day indoor experiment with the TSG100 configured to replicate and retransmit the WWVB signal. The results show that the average synchronization uncertainty of the WWVB frames dropped to 0.499 ms, while availability increased to 100% at every receiver.

6. APPLICATIONS

The results from the previous section suggest that the universal time signal receiver can provide millisecond-level accuracy. Although the synchronization error is relatively high when compared to GPS and synchronization protocols based on message passing (with short re-synchronization intervals), the combination of low cost and ultra-low power consumption makes the universal time signal receiver attractive for applications and deployments that do not require higher than millisecond-level accuracy. Next, we explore some of the representative applications and the potential advantages of utilizing our universal time signal receiver.

Synchronous MAC Protocols. In sensor networks, asynchronous MAC protocols have been preferred over their synchronous counterparts, mainly because synchronous protocols require maintaining synchronized clocks and schedules, adding communication overhead and protocol complexity.

Despite their numerous advantages and widespread use, asynchronous MAC protocols have, due to their random access nature, one inherent weakness when used in a duty cycled network: the sender must keep the radio on while waiting for the intended receiver to wake up. Specifically, sender initiated protocols must continuously transmit (packetized) preambles during the waiting period [3], while receiver initiated protocols must continuously listen for probe packets other motes send when they wake up [11, 43]. On the other hand, a sensor network globally synchronized by hardware devices can improve the communication efficiency of existing synchronous protocols by removing the maintenance overhead caused by message passing synchronization.

Given that MAC protocols in TinyOS tend to operate with millisecond granularity[1] the millisecond level accuracy that the universal time signal receiver provides is sufficient. To illustrate our point, we implemented a time-scheduled version of LPL, which utilizes the time signal receiver to schedule wake up times. Specifically, the sleep interval is divided into slots, with the first and biggest slot reserved for broadcast and multicast traffic, followed by many smaller slots for unicast traffic. We compare the resulting duty cycles by running CTP on top of the standard LPL with a one second sleep interval and our scheduled LPL on a small testbed of TelosB motes. Using a multi-hop network consisting of 10 nodes, with every node generating one packet every 60 seconds, we found the different duty-cycles shown in Figure 9. The chart clearly shows that the duty-cycle improvements are significant, with the average duty-cycling across the 10 motes being reduced from 5.8% to 1.0%. Both protocols achieved equivalent packet delivery ratios ($> 99\%$).

[1]Popular IEEE 802.15.4 radios require multiple milliseconds to transmit and receive packets.

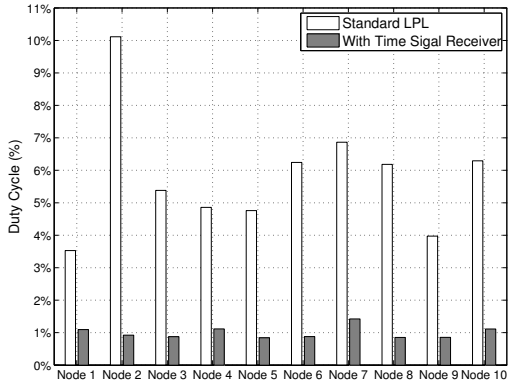

Figure 9: Duty cycles of running CTP on top of standard LPL with a one second sleep interval and modified time signal receiver assisted LPL.

Latency Reduction. Sensor networks can also utilize the universal time signal receivers to improve real time performance. For example, in the case of CTP the MAC layer could be modified in such a way that children nodes are scheduled to transmit prior to their parents, significantly decreasing the end-to-end delay. Our universal time signal receiver could also improve existing low-latency protocols such as RT-Link [38], since their time signal receiver consumed an order of magnitude more power. Replacing it with ours would decrease the power consumption used for synchronization by 88%.

Sparse Networks. One important property of the universal time signal receiver is its independence from radio connectivity among the network's motes. It is therefore well-suited for sparse sensor networks that do not have consistent network connectivity.

For example, some sparse sensor networks employ mobile sinks to collect data [13, 44]. Message passing based synchronization protocols cannot be directly applied in these settings because the sensor network does not form a single connected topology. In this case, GPS receivers and Real Time Clocks (RTC) are the two remaining options.

Depending on the application requirements, GPS can be used when nanosecond accuracy is needed, while, RTCs can be used when lower accuracy is sufficient. High accuracy RTCs can provide $\pm$3.5ppm from -40°C to +85°C, or $\pm$ 2 minutes per year [27].

This is sufficient for applications such as [33], however, for applications that need higher synchronization accuracy, GPS becomes a necessity. Our time signal receiver can provide millisecond level accuracy, and therefore fills the big gap between RTC and GPS.

Drop-in Replacement for GPS. GPS provides the best possible timing accuracy for sensor networks. However, as discussed above, GPS can be an overkill for applications that do not mandate nanosecond or even microsecond accuracy. Some of these deployments end up using GPS because even high accuracy RTCs ([27])

still incur a few minutes drift per year.

For example, Gupchup et al. cited the need for an accuracy in the order of seconds yet still connected a GPS receiver to an anchor mote from which global time was disseminated to the rest of the network by message passing [16]. Considering the low cost and power consumption of our universal time signal receiver, each mote could conceivable be equipped with a time signal receiver easily meeting the time synchronization requirement. Even in the volcano monitoring sensor network deployed by Werner-Allen et al. [46], the desired timing accuracy was only 10 ms, which means our time signal receiver would have been sufficient.

Network-Wide Wakeup. Some sensor networks do not have real time requirements, but rely on network wide wakeup in order to perform data routing. For example, in the Koala data collection system [31], the base station wakes up the whole network and collects the data buffered on each mote. Waking up a duty-cycled network is a special case of network flooding and is taxing on the radio communication budget. On the other hand, with access to global time, waking up a network is equivalent to specifying a rendezvous time, with motes turning on their radio when the rendezvous time arrives. For networks that generate periodic data at a pre-determined rate, the wakeup time can even be pre-programmed.

Failure-Prone Sensor Networks. Motes in failure-prone networks (e.g., networks in which motes are primarily powered by energy harvesting), are expected to experience frequent power resets or even prolonged down times. These events would also cause the mote's local clock to reset. Depending on the desired clock accuracy, motes could use universal time signal receivers to recover the global time after every reset.

7. SUMMARY

We present a mote peripheral that leverages the availability of time signals transmitted by radio stations around the globe to provide access to UTC time with millisecond-level accuracy and $<$ 100 μA current draw. While not as accurate as GPS or message passing protocols such as FTSP, the energy consumption of this peripheral is several orders of magnitude lower than these alternatives, providing sensor network applications an attractive trade-off between accuracy and energy efficiency. We show that this *Universal Time Signal Receiver* can be used in both indoor and outdoor deployments and outline how sensor networks can leverage universal availability to global time that is practically free to further improve their energy efficiency.

With both the WWVB and DCF77 being driven by highly accurate atomic clocks and signals transmitted with microsecond precision, it is disappointing that the accuracy of the CME6005 receiver is only at the millisecond level. Radio chips of similar size, cost, and power consumption, such as the MAS-OY MAS9180 [26] and HKW UE6015 [19], all report similar accuracy. Nevertheless, radios, such as the Meinberg PZF511 [28], that use a different receiver technology are capable of achieving microsecond precision but at a much higher price and power consumption. We intend to investigate the trade-offs between using a more expensive receiver and the increased accuracy that it can achieve.

More information on the prototype universal time signal receiver is available online at `http://hinrg.cs.jhu.edu`.

8. ACKNOWLEDGMENTS

Special thanks to Philippe Bonnet and University of Copenhagen for facilitating the DCF77 experiments and Jan Beutel for shepherding this paper. This material is based upon work partially supported by the National Science Foundation under grants #0834470, #0754782, #0546648.

9. REFERENCES

[1] J. Allison, R. Hormigo, and E. Jovanov. A low-power geographically distributed data acquisition system with wwvb synchronization. In *System Theory, 2003. Proceedings of the 35th Southeastern Symposium on*, pages 162 – 166, mar. 2003.

[2] A. Arora, P. Dutta, S. Bapat, V. Kulathumani, H. Zhang, V. Naik, V. Mittal, H. Cao, M. Demirbas, M. Gouda, Y.-R. Choi, T. Herman, S. S. Kulkarni, U. Arumugam, M. Nesterenko, A. Vora, and M. Miyashita. A line in the sand: A wireless sensor network for target detection, classification, and tracking. *Computer Networks*, pages 605–634, 2004.

[3] M. Buettner, G. Yee, E. Anderson, and R. Han. X-MAC: A Short Preamble MAC Protocol for Duty-Cycled Wireless Sensor Networks. In *Proceedings of the 4^{th} ACM SenSys Conference*, 2006.

[4] N. Burri, P. von Rickenbach, and R. Wattenhofer. Dozer: ultra-low power data gathering in sensor networks. In *Proceedings of the 6^{th} IPSN Conference*, 2007.

[5] C-MAX. Time Signal Generator Shop Floor Booster. Available at `http://www.c-max-time.com`, 2007.

[6] C-MAX. CME6005 Datasheet. Available at `http://www.c-max-time.com`, 2008.

[7] M. Chang. Power efficient duty-cycling with ultra low-power receivers. University of Copenhagen, Master Thesis, 2006.

[8] Y. Chen, R. Tron, A. Terzis, and R. Vidal. Corrective consensus: Converging to the exact average. In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, pages 1221–1228, Dec 2010.

[9] C. Chiasserini and R. Rao. Combining paging with dynamic power management. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 996 –1004 vol.2, 2001.

[10] H. Dai and R. Han. Tsync: a lightweight bidirectional time synchronization service for wireless sensor networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 8(1):125–139, 2004.

[11] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis. Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (Sensys)*, pages 1–14, Nov 2010.

[12] J. E. Elson, L. Girod, and D. Estrin. Fine-grained network time synchronization using reference broadcasts. In *Proceedings of the 5^{th} Symposium on Operating Systems Design and Implementation (OSDI)*, pages 147–163, Dec. 2002.

[13] M. D. Francesco, K. Shah, M. Kumar, and G. Anastasi. An adaptive strategy for energy-efficient data collection in sparse wireless sensor networks. In *EWSN*, pages 322–337, 2010.

[14] L. Gu, D. Jia, P. Vicaire, T. Yan, L. Luo, A. Tirumala, Q. Cao, T. He, J. A. Stankovic, T. Abdelzaher, and B. Krogh. Lightweight Detection and Classification for Wireless Sensor Networks in Realistic Environments. In *Proceedings of the 3^{rd} ACM Conference on Embedded Networked Sensor Systems (SenSys)*, Nov. 2005.

[15] C. Guo, L. C. Zhong, and J. Rabaey. Low power distributed mac for ad hoc sensor radio networks. In *Global*

Telecommunications Conference, 2001. GLOBECOM '01. IEEE, volume 5, pages 2944 –2948 vol.5, 2001.

[16] J. Gupchup, D. Carlson, R. Musaloiu-E., A. Szalay, and A. Terzis. Phoenix: An epidemic approach to time reconstruction. In *Proceedings of the Seventh European Conference on Wireless Sensor Networks (EWSN)*, pages 17–32, Feb. 2010.

[17] J. Y. Halpern and I. Suzuki. Clock synchronization and the power of broadcasting. *Distrib. Comput.*, 5(2):73–82, 1991.

[18] P. Hetzel. Time dissemination via the lf transmitter dcf77 using a pseudo-random phase-shift keying of the carrier. 2nd European Frequency and Time Forum, March 1988.

[19] HKW. UE6015 Datasheet. Available at `http://www.hkw-elektronik.de/pdfenglisch/UE6015%20-%20DE%20V4.1.pdf`, 2006.

[20] P. Huang, H. Oki, Y. Wang, M. Martonosi, L. Peh, and D. Rubenstein. Energy-efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences with ZebraNet. In *Proceedings of the Tenth Internationla Conferece on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, Oct. 2002.

[21] C. Lenzen, P. Sommer, and R. Wattenhofer. Optimal clock synchronization in networks. In *Proceedings of the 7^{th} ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 225–238, New York, NY, USA, 2009. ACM.

[22] M. A. Lombardi. Nist time and frequency services. In *NIST Special Publication 432*, 2002.

[23] R. Mangharam, A. Rowe, R. Rajkumar, and R. Suzuki. Voice over sensor networks. In *Proceedings of the 27^{th} IEEE International Real-Time Systems Symposium (RTSS)*, pages 291–302, Washington, DC, USA, 2006. IEEE Computer Society.

[24] M. Marot, B. Kusy, G. Simon, and A. Ledeczi. The flooding time synchronization protocol. In *Proceedings of the 2^{nd} ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 39–49, Nov. 2004.

[25] C. Marshall, C. Fenger, and P. Gough. Instantaneous, low-power geotagging: Capture & Process is the next killer application for digital photography. Available at: `http://www.u-blox.com/images/downloads/Product_Docs/u-blox_Capture_&_Process_whitepaper_(GPS.CP-X-09000).pdf`, 2010.

[26] MAS-OY. MAS9180 Datasheet. Available at `http://www.mas-oy.com/archive/da9180.pdf`, 2010.

[27] MAXIM. DS3231 Datasheet. Available at `http://www.maxim-ic.com/`, 2010.

[28] Meinberg. PZF511 Datasheet. Available at `http://www.meinberg.de/download/docs/manuals/english/pzf511.pdf`, 2009.

[29] Microchip. PIC16LF1827 Data Sheet. Available at `http://www.microchip.com/`, 2010.

[30] MoteIV Corporation. Tmote Sky. Available at: `http://www.moteiv.com/products/tmotesky.php`.

[31] R. Musăloiu-E., C.-J. Liang, and A. Terzis. Koala: Ultra-low power data retrieval in wireless sensor networks. In *Proceedings of the 7^{th} international symposium on information processing in sensor networks (IPSN)*, pages 421–432, April 2008.

[32] National Institute of Standards and Technology. Available at http://www.nist.gov/pml/div688/grp40/wwvb.cfm.

[33] T. Naumowicz, R. Freeman, H. Kirk, B. Dean, M. Calsyn, A. Liers, A. Braendle, T. Guilford, and J. H. Schiller. Wireless sensor network for habitat monitoring on skomer island. In *SenseApp*, 2010.

[34] L. Niemann, M. Venzke, C. Renner, and V. Turau. Clock Synchronization of TinyOS-based Sensor Networks with DCF77. pages 45–46, Aug. 2009.

[35] Physikalisch-Technische Bundesanstalt. DCF77.

[36] Radio World. NIST Eyes East Coast Version of WWVB. Available at http://www.radioworld.com/article/9286, Jan. 2008.

[37] A. Rowe, V. Gupta, and R. R. Rajkumar. Low-power clock synchronization using electromagnetic energy radiating from AC power lines. In *Proceedings of the 7^{th} ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 211–224, New York, NY, USA, 2009. ACM.

[38] A. Rowe, R. Mangharam, and R. Rajkumar. RT-Link: A Time-Synchronized Link Protocol for Energy- Constrained Multi-hop Wireless Networks. In *Proceedings of the 3^{rd} Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks (SECON)*, volume 2, pages 402 –411, sep. 2006.

[39] T. Schmid, Z. Charbiwala, Z. Anagnostopoulou, M. B. Srivastava, and P. Dutta. A case against routing-integrated time synchronization. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 267–280, 2010.

[40] T. Schmid, R. Shea, Z. Charbiwala, J. Friedman, M. B. Srivastava, and Y. H. Cho. On the interaction of clocks, power, and synchronization in duty-cycled embedded sensor nodes. *ACM Trans. Sen. Netw.*, 7(3):1–19, 2010.

[41] L. Selavo, A. Wood, Q. Cao, A. Srinivasan, H. Liu, T. Sookoor, and J. Stankovic. Luster: Wireless Sensor Network for Environmental Research. In *Proceedings of the 5^{th} ACM Conference on Embedded Networked Sensor Systems (SenSys)*, Nov. 2007.

[42] G. Simon, M. Maroti, A. Ledeczi, G. Balogh, B. Kusy, A. Nadas, G. Pap, J. Sallai, and K. Frampton. Sensor Network-Base Countersniper System. In *Proceedings of SenSys 2004*, Nov. 2004.

[43] Y. Sun, O. Gurewitz, and D. B. Johnson. RI-MAC: a receiver-initiated asynchronous duty cycle mac protocol for dynamic traffic loads in wireless sensor networks. In *Proceedings of the 6^{th} International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 1–14, Nov 2008.

[44] O. Tekdas, V. Isler, J. H. Lim, and A. Terzis. Using mobile robots to harvest data from sensor fields. *Wireless Commun.*, 16(1):22–28, 2009.

[45] Texas Instruments. CC2420: 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver. Available at Available at `http://www.ti.com/lit/gpn/cc2420`, 2006.

[46] G. Werner-Allen, K. Lorincz, J. Johnson, J. Lees, and M. Welsh. Fidelity and yield in a volcano monitoring sensor network. In *Proceedings of the 7^{th} USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Nov 2006.

[47] W. Ye, J. Heidemann, and D. Estrin. An Energy-Efficient MAC Protocol for Wireless Sensor Networks. In *Proceedings of IEEE INFOCOM*, 2002.