
Elapsed time on arrival: a simple and versatile primitive for canonical time synchronisation services

Branislav Kusý

Institute for Software Integrated Systems,
Vanderbilt University,
Nashville, Tennessee 37203, USA
E-mail: branislav.kusy@vanderbilt.edu

Prabal Dutta and Philip Levis

Computer Science Division,
University of California,
Berkeley, California 94720, USA
E-mail: prabal@cs.berkeley.edu E-mail: pal@cs.stanford.edu

Miklós Maróti and Ákos Lédeczi*

Institute for Software Integrated Systems,
Vanderbilt University,
Nashville, Tennessee 37203, USA
E-mail: miklos.maroti@vanderbilt.edu E-mail: akos.ledeczi@vanderbilt.edu
*Corresponding author

David Culler

Computer Science Division,
University of California,
Berkeley, California 94720, USA
E-mail: culler@cs.berkeley.edu

Abstract: Time synchronisation is one of the most important and fundamental middleware services for wireless sensor networks. However, there is an apparent disconnect between existing time synchronisation implementations and the actual needs of current typical sensor network applications. To address this problem, we formulate a set of canonical time synchronisation services distilled from actual applications and propose a set of general application programming interfaces for providing them. We argue that these services can be implemented using a simple time-stamping primitive called *Elapsed Time on Arrival* (ETA) and we provide two such implementations. The Routing Integrated Time Synchronisation (RITS) is an extension of ETA over multiple hops. It is a reactive time synchronisation protocol that can be used to correlate multiple event detections at one or more locations to within microseconds. Rapid Time Synchronisation (RATS) is a proactive timesync protocol that utilises RITS to achieve network-wide synchronisation with microsecond precision and rapid convergence. Our work demonstrates that it is possible to build high-performance timesync services using the simple ETA primitive and suggests that more complex mechanisms may be unnecessary to meet the needs of many real world sensor network applications.

Keywords: sensor networks; ad hoc networks; ubiquitous computing; time synchronisation; clock synchronisation; clock drift; multi-hop; medium access control; packet delay.

Reference to this paper should be made as follows: Kusý, B., Dutta, P., Levis, P., Maróti, M., Lédeczi, A. and Culler, D. (xxxx) 'Elapsed time on arrival: a simple and versatile primitive for canonical time synchronisation services', *Int. J. Ad Hoc and Ubiquitous Computing*, Vol. x, No. x, pp.xxx-xxx.

Biographical notes: Branislav Kusy is a Research Assistant at the Institute for Software Integrated Systems, Vanderbilt University. His current research interests include Time Synchronisation, Radio Based Ranging and Localisation of Sensornets. He is pursuing his PhD in the Electrical Engineering and Computer Science Department at Vanderbilt University, Nashville.

Prabal Dutta is a doctoral student in the Computer Science Division of the Electrical Engineering and Computer Sciences Department at the University of California, Berkeley, where he is a NSF Fellow. His advisor is David Culler. His interests include Sensor Network Design at the hardware-software interface.

Philip Levis was a doctoral student in the Computer Science Division of the Electrical Engineering and Computer Sciences Department at the University of California, Berkeley, at the time of this work, and his advisor was David Culler. He is currently an Assistant Professor of Computer Science at Stanford University. His research interests are Sensor Network Systems, Protocols and Programming. He implemented TOSSIM, the de facto simulator of the sensor network community and helped design nesC, the TinyOS programming language.

Miklos Maroti was a Research Assistant Professor at the Institute of Software Integrated Systems, Vanderbilt University at the time of this work. Currently, he is Assistant Professor at University of Szeged, Hungary. His research interest includes Formal Specification and Analysis of Embedded Systems, and Active Libraries of Middleware Components. He received a PhD in Mathematics from Vanderbilt University.

Akos Ledeczi is a Senior Research Scientist at the Institute for Software Integrated Systems, Vanderbilt University. His current research interests include Model-based Synthesis and Simulation of Embedded Systems. He received his PhD in electrical engineering from Vanderbilt University.

David Culler is a Professor of Computer Science at University of California, Berkeley, a former Director of Intel Research Berkeley, and a member of the National Academy of Engineering. He leads the TinyOS project.

1 Introduction

Time synchronisation has been extensively studied in the context of distributed systems theory, internetworks, local area networks, real-time control systems, wireless sensor networks (WSNs), and many applications. These different contexts have led to different notions of time. A common notion of time allows us to reason about the ordering of events, the causal relationships and correlations between events, the rate of change of observations over time, and also enables the coordination of future actions. Events in distributed systems can be specified in the *temporal*, *causal*, *logical*, and *delivery* order, as well as expressed in *local* and *global* time frames (Kopetz, 1997; Lamport, 1978; Reichenbach, 1957).

In the context of WSNs, timesync commonly refers to the problem of synchronising clocks across a set of sensor nodes which are connected to one another over single- or multi-hop wireless communications channel. Many existing sensor network applications require temporal ordering of events either relative to other events or relative to an absolute external frame of reference. A variety of time synchronisation algorithms have been proposed to address both the temporal ordering of events and the more general problem of establishing a virtual global timebase across a network. Often, these timesync algorithms require nodes to involve third parties (such as reference broadcasters), perform two-way handshakes, exchange multiple messages, or perform computationally expensive optimisations, for example, to maintain synchronised clocks and a consistent sense of network time. Section 2 surveys the related work on time synchronisation in sensors networks, describes how they correlate with the canonical services and what sources of errors they eliminate.

Since all of these operations consume extra system resources, but not all applications have identical synchronisation needs, we may find that system resources are being wasted unnecessarily. This suggests that these more complex and intricate mechanisms may be inappropriate to meet the needs of many common sensor network applications. We suggest that a better match between applications and timesync implementations is possible, but it requires a better understanding of the ways in which applications use timesync services. Section 3 explores the time synchronisation requirements of common sensor applications and distils a set of canonical timesync services from several application domains into a set of timesync application programming interfaces.

Section 4 analyses the sources of time-stamping jitter and physical clock variation. We suggest several mechanisms that exploit the rich interfaces between the software and hardware layers to minimise the non-determinism. Our implementation of these mechanisms is a time-stamping primitive called elapsed time on arrival (ETA), which is a medium access control layer plug-in that reduces or eliminates nearly all sources of time-stamping jitter.

Timesync services can utilise ETA as a low-level primitive provided by the operating system that abstracts underlying hardware and drivers. Porting a timesync service to a different hardware platform requires simply porting the ETA primitive to that platform, promoting loose-coupling and reuse of timesync code. To demonstrate the versatility of ETA, we implemented two general timesync services. Routing integrated time synchronisation (RITS) is a reactive time synchronisation protocol that can be used to correlate multiple event detections at one or more locations to within

microseconds. Rapid Time Synchronisation (RATS) is a proactive timesync protocol that achieves network-wide synchronisation with microsecond precision and rapid convergence. The design details of ETA, RITS, and RATS are presented in Section 5 and a performance evaluation is presented in Section 6. Section 7 presents the conclusions.

Conceptually similar time-stamping primitives have been proposed in the literature but no earlier implementation has achieved comparable performance. As a consequence, other implementations have resorted to more complex designs to mitigate the effects of timing jitter. Our work suggests that these more complex timesync mechanisms may be unnecessary to meet the needs of many common sensor network applications.

2 Related work

WSNs present a unique set of challenges not found in traditional networks such as unreliable, and possibly, unidirectional links, dense and dynamically changing topology and extremely limited energy, computational power and memory resources. Consequently, many results in the rich literature on time synchronisation were not directly applicable in WSNs, and a number of new algorithms and protocols have been proposed.

Sivrikaya and Yener (2004) present a survey of sensor network time synchronisation protocols. Their paper outlines the synchronisation problems using computer clocks as the timebase, identifies common challenges for synchronisation methods, and enumerates the requirements of synchronisation schemes for sensor networks. One of these requirements, called *immediacy*, quantifies the allowable latency when reporting an event to a sink node. The paper claims that immediacy requirements might prevent the use of certain protocols which, the paper claims, in turn, require that nodes be *presynchronised* at all times. We show this claim to be inaccurate; that presynchronisation is not necessary; and that our post-facto approach can convey the event time just as rapidly.

A more in-depth analysis of the timesync problem is presented by Romer et al. (2005) which extends significantly the analysis of essential issues surrounding timesync for sensor networks presented by Elson and Romer (2002). This book chapter describes a system model, identifies classes of synchronisation, presents synchronisation techniques, reviews and classifies ten concrete timesync algorithms from the literature, proposes evaluation strategies, and casts time synchronisation as a special case of the calibration problem. Of particular interest to us are the descriptions, classifications, and evaluations of the ten timesync protocols: Time-stamp Synchronisation (TSS) (Romer, 2001), Reference Broadcast Synchronisation (RBS) (Elson, 2002), Tiny-Sync and Mini-Sync (TS/MS) (Sichitiu and Veerarittiphan, 2003), Lightweight Time Synchronisation (LTS) (van Greunen and Rabaey, 2003), Timing-Sync Protocol for Sensor Networks (TPSN) (Ganerwal et al., 2003), TSync (Dai and Han, 1994),

Interval-Based Synchronisation (IBS) (Marzullo and Owicki, 1983; Blum et al., 2004; Meier et al., 2004; Schmid and Schossmaier, 1997), Flooding Time Synchronisation Protocol (FTSP) (Maróti et al., 2004), Asynchronous Diffusion (AD) (Li and Rus, 2004), and Time Diffusion Synchronisation (TDP) (Su and Akyildiz, 2004). Most relevant to our work are TSS, RBS, TPSN, FTSP, and the timesync protocol used in a structural monitoring application by Xu et al. (2004).

The Time-Stamp Synchronisation (TSS) approach is conceptually similar to our time-stamping primitive. One drawback to the TSS approach is the requirement for multiple floating-point operations at each node which, on mote-class devices, is expensive. Another, more significant drawback stems from the assumption that it is difficult to accurately estimate message delays between neighbouring nodes. This assumption is due to the queuing, randomness, and contention that typically occur in the media access control layer of many networks on desktop-class computers. However, with the advent of specially-designed operating systems such as TinyOS (Hill, 2000) and integrated sensor nodes like the Mica motes, event callbacks that are invoked at the *actual time of message transmission* are possible.

The Reference Broadcast System (RBS), perhaps the best known timesync algorithm in the sensor network regime, uses a transmitter node to synchronise the clocks of two receiver nodes to each other. The RBS approach time-stamps messages only on the receiver side, which eliminates the delays on the sender side, most notably in the medium access control layer. The accuracy of the RBS time-stamping reported by the authors is approximately 11 μ s on the Mica platform. Least square linear regression is used to account for the clock drifts which results in 7.4 μ s average error between two motes after a 60 second interval. The RBS solution for the multi-hop scenario uses concepts similar to one of our proposed canonical services. However, because of the extra coordination overhead of reference broadcasters, the implementation is considerably more complicated and less energy efficient than our proposed solution in Section 5.2.

The Timing-Sync Protocol for Sensor Networks (TPSN) eliminates sources of time-stamping errors by making use of the implicit acknowledgments to transmit information back to the sender. This protocol gains additional accuracy over RBS due to time-stamping of the radio message twice and averaging these time-stamps. Since TPSN relies on a two-way information exchange, messages cannot be broadcasted, which results in higher communication load. The authors of TPSN algorithm implemented both TPSN and RBS on the Mica platform using a 4 MHz clock for time-stamping, and compared the precision of the two algorithms. The resulting average errors for a single hop case for two nodes are 16.9 μ s and 29.1 μ s for the TPSN and RBS algorithms, respectively.

The Flooding Time Synchronisation Protocol (FTSP) improves on the ideas in RBS and TPSN. It is a proactive protocol in which every node periodically broadcasts timesync messages. The local clock of a single dynamically

elected node provides the global time in the network. A time synchronisation hierarchy rooted at this node is created in which nodes are synchronised by receiving global time estimates from nodes closer to the root. The applied flood-based communication protocol in FTSP provides a very robust network, and still incurs only small network traffic. The algorithm uses a fine-grained clock, MAC-layer time-stamping with several jitter reducing techniques and clock drift estimation to achieve high precision. On the Mica2 platform the reported average error of the algorithm is 1.48 μs for the single hop case, and 0.5 μs per hop in a 6-hop network for the multi-hop case. However, the initial convergence time of FTSP on a 60-node 6-hop network is 10 minutes.

The FTSP time-stamping reduces non-determinism to such a degree that RBS's receiver-receiver and TPSN's two-way sender-receiver synchronisation become unnecessary, and one-way sender-receiver synchronisation suffices. FTSP accomplishes this by identifying, estimating, and compensating for almost all sources of uncertainties in the radio message delivery. We fully leverage and extend FTSP's results to build our time-stamping primitive presented in detail in Section 5.

Hardware accelerators represent the most promising approach to reducing jitter uncertainty to the level of the system clock granularity. Hill and Culler (2002) observed that unlike wide-area time synchronisation protocols such as NTP, in sensor networks we may be able to determine all sources of communication delay. By exposing all sources of delay to the application, they were able to synchronise a pair of nodes to within 2 μs of each other using sender-receiver time-stamping. The sources of jitter included several factors: raw RF transmission jitter of $\pm 1 \mu\text{s}$ on the transmission propagation due to the internals of the radio, hardware capture of the arriving pulse to within an accuracy of $\pm 0.25 \mu\text{s}$, and synchronisation code path delay of an additional $\pm 0.625 \mu\text{s}$. This implementation was only possible because of the rich interface between the radio hardware and system software.

Dutta (2004) advocates the use of reactive timesync algorithms for random event detection, arguing that reduced power consumption and low latency due to the lower messaging frequency and no need for calibration, respectively, map well to the needs of event detection. Xu et al. (2004) used an approach similar to ours for time-stamping data in a WSN for structural monitoring. Both of these works focused on specific applications and neither generalised their approach to be broadly applicable.

3 Canonical services

Many sensor network applications require time synchronisation but the *model* of synchronisation that is needed varies greatly. In practice, many application developers resort to the virtual global time service without considering other alternatives. Virtual global time is chosen because it is conceptually simple, but since it is proactive and message-intensive, virtual global time may use system

resources unnecessarily. We argue that abstracting common timesync usage patterns from a number of existing applications and standardising the interfaces of these abstractions will help developers better identify the synchronisation needs of sensor network applications and help them choose the most efficient model for their needs.

This section presents a set of canonical time synchronisation services derived from several well-known and common sensor network applications. Our focus is on the use of the services in *actual applications*, so we also propose a set of application programming interfaces (APIs) that could be used by these or other similar applications to access timesync services. However, this list is neither complete nor predictive; we did not abstract every single use of timesync found in the literature nor did we attempt to identify future models of synchronisation. Nevertheless, this list of timesync services, each of which can be implemented using the proposed ETA primitive, suggests that other protocols, not on the list, could be implemented using ETA as well.

Table 1 summarises the canonical services, APIs and typical applications. Interestingly, we tried hard to find an application that truly required global time synchronisation. This indicates a possible mismatch between the areas which have garnered much of the community's attention and the actual needs of common WSN applications.

Table 1 Canonical time synchronisation services, their typical applications, and proposed API calls and events that can be used in a general implementation. Events are callbacks from a service layer to a higher layer and are of the form OnXXX

<i>Canonical service</i>	<i>Typical applications</i>	<i>APIs</i>
Event time-stamping	Intrusion detection	SendToSink
	Countersniper system	SendToAll
	Source localisation	OnReceive
Data series time-stamping	Habitat monitoring	SetAccuracy
	Structural monitoring	MarkStart
	Environment monitoring	MarkStop
	Volcano event monitoring	
	Target classification	
Virtual global time	Beamforming	
	Debugging traces	GetGlobalTime
		LocalToGlobal
Coordinated action		GlobalToLocal
	Scheduled data collection	Schedule
	Communications scheduling	OnSchedule

The remainder of this section delves into the canonical services and their underlying abstractions, APIs, and motivating applications in greater detail. For the remainder of this paper, we restrict the discussion to only those

applications which require temporal order, since logical and delivery orders are trivially implemented with messaging and causal ordering requires a temporal ordering capability.

3.1 Event time-stamping

A single, isolated event is observed and time-stamped by one or more nodes at possibly different times. The observation from each node is to be collected at one or more sink nodes, either individually or in aggregate. The sink node must determine the time-stamp of each event in the sink's local time, *regardless of whether the sink is coordinated with an external source of time like GPS*. This service extends to multiple events through repeated invocation.

An event time-stamping API that consists of a `Packet` data structure augmented with an `eventTime` field, along with calls to `convergecast` a packet (i.e., send a packet toward the sink node)

```
command bool SendToSink(Packet packet)
```

flood a packet (network-wide) to all other nodes

```
command bool SendToAll(Packet packet)
```

and signal an event when new packet is available at a node (which is signalled at every intermediate hop)

```
event void OnReceive(Packet *packet)
```

appears sufficient to meet the `timesync` need of these applications.

3.2 Data series time-stamping

One or more observer nodes are periodically sampling some phenomena. The data series from each node are to be collected at a sink node for correlated analysis. Data series time-stamping requires time-stamps on each sample that are both accurate with regard to a series (intra-series accuracy) as well as accurate across series (inter-series accuracy). Unlike single event time-stamping, which is a discrete event and can benefit from discrete corrections such as a single RBS exchange, data series time-stamping is a continuous effort that requires continuous maintenance.

The precision of the data series time-stamps varies across applications. The phenomena in habitat monitoring are at a slow enough time scale that delivery order at a base station is a reasonable form of time-stamping. In contrast, the phenomena in structural health monitoring require in-network stamping to within a fraction of a packet transmission time. At the extreme are the applications in which the precise phase difference of a signal arriving at multiple nodes is important. For example, acoustic source localisation and beamforming applications require that a signal phase be measured to within a fraction of a cycle, which for an acoustic signal in the kHz range might be just a few tens of microseconds. In many signal processing algorithms; samples cannot be dropped to compensate for clock errors. The challenge is to compare observations when

the sampling frequencies of the sources are not identical, the observations are collected over a (potentially long) period of time, and samples cannot be dropped or added (i.e., time must not have any sudden jumps).

A data series time-stamping API that includes calls to set the required inter-series accuracy

```
command bool SetAccuracy(int micros)
```

indicates the beginning of a data series (which initiates any under-the-hood work needed to meet the specified accuracy)

```
command bool MarkStart()
```

and indicates the end of a data series (which ends any under-the-hood work started to meet the specified accuracy)

```
command bool MarkStop()
```

could meet the `timesync` need of these applications provided that the underlying implementation could meet the constraints. Each `timesync` implementation might have its own method for satisfying the accuracy requirements of the application and no implementation might be able to satisfy all constraints.

3.3 Virtual global time

Virtual global time is perhaps conceptually the simplest of the time synchronisation services to understand but is often the one which requires the most overhead to achieve. The basic idea is simple: a single, virtual, and globally shared timebase is accessible at each node in the network. An important question is where the timebase originates. Options include an external source like GPS, a distinguished root node, or some kind of network statistic as the network mean.

A virtual global time API that includes calls to read the current global time

```
command bool GetGlobalTime(int *time)
```

and to convert between the local times and their corresponding global times

```
command bool LocalToGlobal(int *time)
```

```
command bool GlobalToLocal(int *time)
```

appears sufficient to meet the `timesync` need of these applications.

3.4 Coordinated action

Coordinated action is similar to event time-stamping, except that the event is in the future. Applications that require coordinated action include scheduled data collection and structural health monitoring. In the former, the action is a discrete data collection; depending on the different required fidelities, the synchronisation needed for coordinated action can be used for the events themselves. In the latter, the action is a data time series; although the coordination can satisfy the inter-node accuracy, the data series service must still be responsible for intra-node accuracy.

The challenge in coordinated action is that, unlike simple time-stamping, it cannot benefit from post-facto synchronisation such as an RBS exchange. However, unlike virtual global time, the nodes do not need a continuous synchronised time value: they merely need to agree on a single point in time. While an implementation of this service can depend on virtual global time, other implementation options exist.

We propose a coordinated action API that includes calls to schedule an action, in a node's *local* time, with optional quality of service parameters and number of repetitions

```
Command bool Schedule(int *time)
Command bool SCHEDULE(int *time, int.
accuracy)
Command bool SCHEDULE(int *time, int
accuracy, int. repeat)
```

and signal an event when it is time to act, for the *count*-th time, on some previously received coordinated action request

```
event void OnSchedule(int count).
```

4 Sources of time synchronisation error

This section establishes the terminology and identifies the various sources of error that contribute to time synchronisation uncertainty. We analyse two fundamentally different types of errors:

- time-stamping jitters encountered when transmitting radio messages between nodes
- time errors caused by the manufacturing and operating differences of physical clocks.

We also provide guidelines on how to design and implement high precision one-way sender-receiver time-stamping service for broadcasted messages.

4.1 Time-stamping jitter

The sources of errors in message time-stamping need to be carefully analysed and compensated for, because non-deterministic delays in radio message delivery have magnitudes larger than the required precision of time synchronisation. We use the following decomposition of these errors, first introduced by Kopetz and Ochsenreiter (1987) and later extended by Kopetz and Schwabl (1989), Horauer et al. (2002) and Ganerival et al. (2003). The presented list is a generalisation of a similar list presented by Maroti et al. (2004).

- *Send time*. The time used to assemble the message and issue the send request to the MAC layer on the transmitter side. This depends on the system call overhead and processor load.

- *Access time*. The delay incurred waiting for access to the transmit channel up to the point when transmission begins. This depends on the channel contention.
- *Transmission time*. The time required for the sender to transmit the message. This depends on the length of the message and the speed of transmission.
- *Propagation time*. The time required for the message to propagate from the sender to the receiver once it has left the sender. This depends only on the distance between the two nodes.
- *Reception time*. The time required for the receiver to receive the message. The transmission and reception times have similar characteristics and overlap, if the propagation time is negligible.
- *Receive time*. Time to process the incoming message and to notify the receiver application. Its characteristics are similar to that of the send time.

We further analyse the sources of uncertainties in the overlapping transmission and reception times by observing an idealised point of radio message, such as the end of a particular byte. We follow the transmission of this idealised point through the software, hardware and physical layers of the wireless channel from the sender to the receiver. The following delays in the propagation of the idealised point seem to be the most important.

- *Interrupt handling time*. The delay between the radio chip raising, and the microcontroller responding, to an interrupt by recording a time-stamp.
- *Encoding time*. The time required for the radio chip to encode and transform a part of the message to electromagnetic waves. This starts when the radio chip initiates the transfer of the idealised point.
- *Decoding time*. The time required for the radio chip on the receiver side to transform and decode the message from electromagnetic waves to binary data. This ends when the radio chip raises an interrupt indicating the reception of the idealised point.

Some radio chips cannot capture the byte alignment of the transmitted message stream on the receiver side and the radio stack has to determine the bit offset of the message from the alignment of a known synchronisation byte. Since the transmission time of a byte is a few hundred microseconds at 38.4 kbps, the delay caused by the incorrect byte alignment must be compensated for.

- *Byte alignment error*. The delay incurred because of the different byte or data segment alignment of the sender and receiver. This is deterministic and can be computed at the receiver side from the bit offset and the speed of the radio.

The list provided above characterises phenomena relevant to the time-stamping of messages transmitted over the wireless channel. We direct the reader to a paper by Maróti et al. (2004) for a more comprehensive analysis of these sources on the Mica2 platform. Note that some radio technologies are not impacted by all these sources of error.

4.2 Physical clocks

Computer systems typically measure time in discrete steps by counting the oscillations of a physical clock. Such clocks are driven by quartz crystals, ceramic resonators, or resistor-capacitor oscillator circuits, depending on the accuracy, stability, power, startup time, and cost requirements of the applications. More advanced solutions are available as ultra-stable oscillators constructed from temperature-compensated sapphire resonators. However, for the majority of applications, common quartz crystals provide a reasonably accurate and cost-effective solution. We assume that a quartz crystal is the basis of the physical clock for the remainder of this section and use the terms crystal and physical clock interchangeably.

A quartz crystal oscillator is an analog component which outputs a continuous-time sinusoidal signal converted to a digital signal by microcontroller. The period of the physical clock sets the length of the discrete time quantum, or *granularity*, of the timebase. All synchronisation protocols encounter errors caused by the insufficient precision of the time representation when an analog-digital boundary is crossed (like event detection or time-stamping). The exception to this is the time-triggered technology that can phase synchronise the clocks across the distributed system:

- *Time representation error*: the delay between an event and the nearest (or next) representable time value

Under ideal circumstances, physical clocks oscillate at a constant frequency. In the real world, manufacturing variations and exposure to out-of-tolerance conditions (e.g., mechanical shock) result in permanent frequency errors of crystals, and variations in temperature and age result in short-term errors of the crystals.

- *Frequency skew*. Under specified conditions at a nominal ambient temperature, the difference between the crystal frequency and the nominal frequency. Usually expressed as $\Delta f/f$ with typical values around ± 50 ppm.
- *Temperature characteristics*. Taking the frequency at a nominal ambient temperature as the reference (usually $+25^\circ\text{C}$), the change in frequency with respect to a change in the ambient temperature. Different crystal types have different temperature characteristics:
 - for tuning-fork crystals, $\Delta f/f$ depends quadratically on the difference between the ambient and the nominal temperature

- for AT-cut crystals, the relationship between $\Delta f/f$ and the temperature difference is described by a 3rd order polynomial equation
- *Ageing*. Amount of frequency drift when operated under the specified conditions for a specified term. A typical value for ageing might be -5 ppm in the first year.

4.3 Minimising the time-stamping jitter

The uncertainties of the send, receive, access and byte alignment times are best eliminated by time-stamping the message in the MAC-layer, as is done in several time synchronisation protocols. The propagation time cannot be calculated or compensated for within the transmission of a single message. However, in static networks it has no jitter and over short distances (less than 300 meters) its duration is negligible (less than one microsecond). We argue that the message needs to be time-stamped at an idealised point, practically when an interrupt is raised by the radio. This eliminates the effect of the overlapping transmission and reception times. The jitter in the interrupt handling time is best eliminated by utilising a capture register on the microcontroller. If the time-stamp is taken in software, it is important to minimise the length of all atomic sections in the application to minimise the amount of time that interrupts are disabled. The encoding time lasts several hundreds of microseconds but usually has very low jitter. On the other hand, the radio technology, signal strength fluctuations and bit synchronisation errors will introduce jitter in the decoding time. The following novel technique can be used to reduce the jitter of the encoding, decoding and interrupt handling times. We record multiple time-stamps both at the sender and receiver sides as the message is being transmitted, then using statistical analysis we arrive at the final time-stamps (one at the sender and one at the receive side) that has lower jitter than the individual ones. Details of the statistical analysis are presented by Maróti et al. (2004).

The latest generation of radios make the processor-radio interface richer without burdening the processor with the overhead of directly coordinating the transmission. Such improvements also enable fine-grained timing information to be communicated from the radio to the processor while using an abstract and high-level interface. For example, the Chipcon CC2420 radio, which supports the IEEE 802.15.4 standard, allows random access into the transmit FIFO during transmission (Chipcon, 2003). Such access to the transmit queue, when coupled with the Start of Frame Delimiter (SFD) output signal – a deterministic hardware interrupt on the CC2420 – enables fine-grained message time-stamping with virtually no jitter uncertainty on the transmit end and ± 0.125 μs of uncertainty on the receive

end. By connecting the SFD signal to a timer capture pin on the processor, accurate time-stamping on the order of the time representation error is possible.

The presented techniques reduce only the jitter of these times, but do not estimate their absolute value. The sum of these delays is the difference between the sender and receiver side time-stamps. This sum can be measured with a one-time calibration procedure utilising a many-to-many message handshake similar to the RBS algorithm. We point out that using software tools only, it is impossible to measure the absolute value of the encoding and decoding times separately, only their sum.

We direct the reader to the paper by Maróti et al. (2004) for a more comprehensive analysis of these techniques on the Mica2 platform in which 1.4 microsecond time-stamping precision was reported. Our key observations is that with deterministic MAC-layer time-stamping, we can eliminate the runtime calibration phase that is required in the work of Romer (2004) and Elson (2002), the need for a ‘third party’ node as suggested by Elson (2002) and the Single-Pulse Synchronisation presented by Elson (2003) for many of the existing applications of sensor networks.

4.4 *The choice of local time*

The first key design choice in any time synchronisation protocol is the source of times that we want to synchronise. We argue that each node must have a free running local clock whose offset and skew is only measured and compensated for, but never altered. This local clock gives a natural time line where events are recorded, actions are scheduled and calculations are performed.

The time representation error of events can never be eliminated, but it can be reduced by using a higher frequency crystal. For example, the Mica2 mote includes a 32.768 kHz crystal which provides 30.52 μ s granularity which is sufficient for many applications. If a finer granularity clock is needed, the Mica2 also includes a 7.3728 MHz crystal which provides 0.1356 μ s granularity.

The actual frequency of a particular crystal instance can take any value within the range specified by the tolerance precision. If a particular crystal’s frequency is measured at multiple temperatures, we can create a calibration table that provides a frequency skew factor for every ambient temperature. Typically, this kind of calibration is performed at manufacture time for instrument-grade electronics like oscilloscopes and counters. The calibration must be repeated periodically because ageing and shock can cause frequency variations. For WSNs that contain hundreds or thousands of nodes, repeated manual calibration is simply not an option. Instead, it might be possible to perform distributed parameter calibration, using the fact that sensor nodes have multiple crystals with different temperature characteristic curves. For example, on the Mica2 motes, the 32 kHz clock is based on a tuning fork crystal while the 7 MHz clock is an AT-cut crystal.

5 **Implementation of canonical services**

We argue that both the development and use of timesync services will be greatly enhanced by shielding developers from the myriad sources of time-stamping errors presented in Section 4 by encapsulating these errors in a timesync primitive. Our reasons are as follow.

- The implementation of a timesync primitive is best done along with the analysis of the sources of error by the hardware domain experts. Consequently, the primitive can provide fine-tuned performance by reducing or eliminating nearly all sources of timesync errors.
- The timesync primitive is well suited to the component based architecture of distributed applications and can simplify the implementation of timesync services.
- The timesync primitive improves the portability of timesync services to a different hardware platform by isolating code changes to the timesync primitive.
- Finally, the timesync primitive can be implemented as a service provided by the operating system, making it easier to access hardware interfaces directly and enable higher precision time synchronisation protocols to be implemented (Ganerwal et al., 2003; Maróti et al., 2004) than would be possible otherwise.

In the remaining part of this section, we introduce our time-stamping primitive and show how it can be used to implement two of the canonical services described in Section 3. We demonstrate the significant decrease in complexity of the implementation of these services when compared to the existing protocols found in the literature.

5.1 *The elapsed time on arrival time-stamping primitive*

Conceptually, the elapsed time on arrival (ETA) algorithm is useful in the case when a sensor node detects a certain phenomenon of interest, or *an event*, and a neighbour of the sensor node needs to know the time of this event. ETA allows neighbouring nodes to establish a common time base by sending a single radio message transmission. In our approach, the times of all events of interest are recorded in the local time of a node and inter-node time conversions occur whenever messages are transmitted from one node to another. At transmission time, the elapsed time since the event occurrence is computed and included as a separate field in the message. On receipt, a node computes the event time by subtracting the elapsed time from the receiving node’s local time.

The implementation details for reducing timesync uncertainty for the Mica2 platform were discussed in Section 4.3 and can be also found in the work by Maróti et al. (2004). The pseudo code of the ETA algorithm is shown in Figure 1.

Figure 1 The elapsed time on arrival (ETA) algorithm

```

struct TimeStampedMessage {
    char[] payload;
    int eventTime;      // elapsed time (msg in transit)
                      // local time of event (otherwise)
    int transmissionTime; // not actually transmitted
}

OnBeginTransmission(TimeStampedMessage msg) {
    msg.transmissionTime = GetLocalTime();
    msg.eventTime = msg.transmissionTime - msg.eventTime;
}

OnBeginReception(TimeStampedMessage msg) {
    msg.receptionTime = GetLocalTime();
}

OnEndReception(TimeStampedMessage msg) {
    msg.eventTime = msg.receptionTime - msg.elapsedTime;
}

```

Each node accesses its local clock or timer through the `GetLocalTime` call. On an event occurrence, the local time is recorded in the `eventTime` field and the event message is enqueued for the transmission. When the event message is scheduled and the idealised point of the message is being transmitted (ostensibly after some delay in the media access control layer), both the sender and the receivers record their current local time into the `transmissionTime` field. The sender also updates the `eventTime` field to contain the elapsed time since the event occurrence. Note that the `transmissionTime` is not transmitted over the wireless channel to the neighbours. After the whole message has been received, the `eventTime` is reconstructed on the receiver side by subtracting the elapsed time from the `receptionTime`. The key enabler of ETA is the ability to update the content of the radio message (to update the `eventTime` field) while the message is being transmitted.

The proposed time-stamping protocol eliminates or reduces all sources of time-stamping error except for propagation time and the effects of the clock skew between the sender and receiver while a message is waiting in the queue. Neither of these errors can be estimated with a single radio message and hence this estimation is left for higher layers to implement. The overhead of the algorithm is a single field in the transmitted message. After message transmission is complete, the sender and all receivers know the time of the transmission, as well as the time of the event in each receiver's local time base.

5.2 Event time-stamping service

The Routing Integrated Time Synchronisation (RITS), a reactive time synchronisation protocol, can be used to obtain event detections at multiple observers in terms of local time of the sink to within microseconds. RITS is an extension of time-stamping primitive ETA over multiple hops. It is achieved by integrating ETA into the multi-hop routing engine DFRF presented by Maroti (2004).

After detecting an event, the observer stores the relevant data in a local data structure along with the local time of the

event occurrence and uses one of the methods of RITS to send the data to the sink. RITS, being a routing engine, sends the packet to the sink along a multi-hop path while maintaining the event time using ETA. At each hop it converts the event time from the local time of the sender to the local time of the receiver.

The details of our implementation are shown in Figure 2. Upon receiving a request to send a data packet, RITS creates a `TimeStampedRITSMMessage` and off-loads all the local time conversion details to the ETA algorithm. Upon receiving a `TimeStampedRITSMMessage`, RITS creates a new packet and associates with it, `message.eventTime` which is the event time converted by ETA to the local time of the receiver. RITS then signals the reception of packet to the application level and possibly enqueues the packet for the further transmission.

Figure 2 Pseudo code of the RITS implementation

```

struct RITS_Packet {
    char[] data;
    int eventTime;      // not actually transmitted
}

SendToSink(RITS_Packet packet) {
    TimeStampedRITSMMessage msg;
    msg.payload = packet.data;
    msg.eventTime = packet.eventTime;
    MessageSend(msg);    // transmit over the radio
}

OnMessageReceive(TimeStampedRITSMMessage msg) {
    RITS_Packet packet = new RITS_Packet();
    packet.data = msg.payload;
    packet.eventTime = msg.eventTime;
    signal OnReceive(&packet);
    Enqueue(packet);    // transmit via routing
}

```

5.3 Virtual global time service

Rapid Time Synchronisation (RATS), a proactive time synchronisation protocol, achieves network-wide timesync in medium size networks with microsecond precision and rapid convergence. The basic idea is the same as in RITS, except that the direction of messages is reversed: instead of convergecast from multiple nodes to the sink, we use a broadcast from a single node, the root, to all other nodes in the network.

Time synchronisation is initiated by the root by broadcasting a message containing two fields, `rootEventTime` and `eventTime`, which are both initialised to the same value in the root's local time (the current time). Network wide broadcast is used to transmit the root's message to all nodes in the network. The `rootEventTime` field is not modified during the broadcast, while ETA converts the `eventTime` field to the local time of the receiver of the message. Thus when a node receives the message, it obtains two times corresponding to the same time instant: `rootEventTime`, in the local time of the root, and `eventTime`, in the local time of the node.

We call the pair $(\text{rootEventTime}, \text{eventTime})$ a synchronisation point and use several of them to estimate the offset and skew of the local clock of receiver to the root's clock. Consequently, the local time of the root becomes the global time in the network.

The most widely used method to estimate the clock drift of a node to the clock of the root is linear regression (Maroti et al., 2004; Elson, 2002): a fixed number of synchronisation points are locally stored in a table at each node, a linear regression line is fit to these points and the slope and intercept of the line is used to estimate the root's time in the future. Since the drift of the clock may change over time (see Section 4), the root needs to rebroadcast synchronisation messages periodically. We use a sequence number field in the RATS message to distinguish new synchronisation messages from old ones. We used RITS to implement RATS (see Figure 3) but it is straightforward to implement RATS directly over ETA.

Figure 3 Pseudo code of the RATS implementation

```

struct RITS_Packet {
    int    rootEventTime;
    int    eventTime;
}
OnRootSend() {
    RITS_Packet packet;
    packet.eventTime = GetLocalTime();
    packet.rootEventTime = packet.eventTime;
    RITS.SendToAll(packet);
}

event RITS.OnReceive(RITS_Packet *packet) {
    LinRegression (packet->rootEventTime, packet->eventTime);
}

```


6 Evaluation

We implemented the ETA, RITS and RATS protocols in TinyOS (Hill, 2000) on the UC Berkeley Mica2 platform. Even though the RITS and RATS protocols are simple extensions of ETA, our experimental evaluation shows that their performance is comparable to, if not better than, the leading timesync algorithms found in the WSN literature.

6.1 Evaluation of RITS: routing integrated time synchronisation

We tested the performance of RITS in an experimental setup using 45 Mica2 motes arranged in a grid. Each node was able to communicate only with its neighbours as shown in Figure 4, forming a 10-hop network. This constraint was enforced in software. The test scenario involved multiple observers detecting the same event and reporting the event detection times to the sink. We simulated the event detection by a radio message that was transmitted by a single node and received by all the nodes that were, at most, two hops away from the sender. Figure 4 shows the sink, the nodes, the radio links and the detection radius of an event.


Figure 4 The layout of the RITS experiment: a node communicates with at most eight nodes and an event is detected at most at 12 nodes



Each node recorded the time of arrival of the message as an event detection time (`message.receptionTime` in Figure 1), created a detection message consisting of the id of the node (`message.nodeID`) and the detection time (`message.eventTime`) and sent this message to the sink. Consequently, the sink received detection messages from up to 12 different nodes. Each such message contained the time of the simulated event at a particular node converted by RITS to the local time of the sink. The times reported by two different nodes should be the same if we neglect the propagation time of a radio message. However, the time-stamping error as well as the clock skew error cause a small variance of the received times. For each group n reported times we calculate the maximum and the average of the $\binom{n}{2}$ pairwise errors in the group. We call these values and average and *maximum* synchronisation error, respectively.

We broadcasted the event simulation message five times with 100 ms delays in a single simulation round. The sender of the event simulation message was chosen randomly both within and across rounds. We initiated consequent simulation rounds with a 30-second period, let the experiment run for one and one-half hours and collected data for 900 simulated events. We plot the histogram of *maximum* synchronisation error in Figure 5. The maximum and average errors of the maximum synchronisation error over all rounds were 80.19 μs and 7.86 μs respectively.

Figure 5 Shows the histogram of the *maximum* pairwise error in the RITS experiment



The performance of RITS is principally affected by two factors. First, variations in clock crystal frequencies – we have observed delays of up to 17 seconds between the event happening and the event notification message arriving to the sink in a particular round. The frequency error of a typical oscillator is in the range of 10–100 ppm, therefore, a 7 MHz MICA2 oscillator can introduce significant errors over these long delays. Second, variations in the jitter of time-stamping – this affects RITS the same way as any other timesync protocol that uses radio message time-stamping to establish a common time base.

The timesync accuracy of RITS as shown by our experiment is slightly worse than the accuracy of published proactive algorithms by Ganeriwal et al. (2003), Maroti et al. (2004) and Elson (2002). However, we argue that this accuracy is sufficient for most applications and that the energy-quality tradeoffs of RITS vs. proactive protocols is acceptable, and even desirable, in many cases.

The benefits of RITS over proactive timesync protocols are as follows. First, the drifts of the local clocks of the nodes do not need to be calculated and maintained in RITS whereas proactive algorithms require periodic resynchronisation. Second, RITS can be powered down completely and woken up by the event itself (as in Lédeczi et al., 2005), which enables more sophisticated power management techniques than proactive protocols which require the nodes to be continuously active, even between synchronisation events, since at least local clocks need to be running. Third, RITS introduces minimal communication overhead; there is no need for time synchronisation related messages as in proactive algorithms and only a single field needs to be included in the transmitted message.

6.2 Evaluation of RATS: rapid time synchronisation


We evaluated the performance of RATS using the same test setup as was used for FTSP (Maróti et al., 2004). We set up

60 Mica2 motes in a 5×12 grid so that each node had up to eight neighbours enforced in software. This is similar to the RITS test scenario shown in Figure 4. The test setup involves 60 synchronising nodes, the reference broadcaster that queries the global time estimates of the synchronising nodes, and the base station that collects the reported global times. The reference broadcaster and the base station are used just for evaluation purposes; they play no role in the synchronisation algorithm.

The root was programmed to transmit synchronisation messages first for 10 seconds with a period of two seconds and then for the rest of the experiment with a period of 30 seconds. The reference broadcaster queried the global time from all the nodes with a period of five seconds during the first two minutes and with a period of 23 seconds (which is a relative prime to 30) the rest of the time. In each reference broadcast we obtained up to 60 reported global times, one of them being the root's time. For each reference broadcast, we computed the *maximum* and the *average* absolute errors as the maximum and the average absolute difference between the root's time and the other times for each reference broadcast round.

We ran the experiment for six hours and achieved network-wide synchronisation (the point in time when at least two synchronisation points had been received by all nodes) only four seconds after switching on the root. Over all rounds, the maximum and average errors of RATS were 26 μ s and 2.7 μ s, respectively. This is in contrast to the 10-minute convergence time, 14 μ s maximum error and 2.3 μ s average error of FTSP (Maroti et al., 2004). We note that the FTSP experiment had only six hops whereas our experiment had 11 hops, so we suspect that the maximum and average errors of RATS and FTSP might have been closer if FTSP were synchronising over 11 hops. Figure 6 shows the maximum and average errors for each round of the experiment.

Figure 6 The RATS experiment showing the maximum and average errors of the reported global times. The figure on the left shows the first 10 minutes and the figure on the right shows the whole experiment. The synchronisation was achieved after 4 seconds and the average error was 2.7 μ s in the 11 hop network



At the beginning of experiment the errors were generally larger than later on. This is expected as initially only a few datapoints were available in the regression table and the clock skew estimate was not as accurate. Moreover, errors

in individual measurements have a much larger effect on the accuracy of the skew estimation, because initially the regression table covers only a few seconds. The data also illustrate this point: when the protocol switched to the

30-second period, the maximum error gradually increased to 26 μs and then gradually decreased as the skew estimates improved. More generally, the tradeoffs between RATS and FTSP are as follows.

Network-wide convergence time of timesync in RATS was shown to be four seconds for 11 hops; it took 10 minutes to achieve network-wide synchronisation of FTSP for a 6-hop network as reported by Maróti et al. (2004). Depending on the accuracy needs of the application it is possible to achieve sub-second convergence using RATS by decreasing the initial synchronisation period even further.

The robustness of RATS is superior to FTSP in two regards. First, the synchronisation error of a node in RATS does not depend on the current clock drift estimation of the nodes on the route from the root but it does for FTSP. Second, the synchronisation message from the root arrives at each node along multiple routes, giving the node multiple local times for each global time. RATS takes the *median* of the received local times as its synchronisation point, which is robust to outliers, and allows RATS to tolerate faulty or adversarial nodes whereas FTSP cannot.

Unlike FTSP, RATS cannot maintain network-wide synchronisation if the network is partitioned and is vulnerable to a root failure. However, since most WSN applications depend on a gateway for wide-area network access, we have the base station already. Since RATS floods the network with radio messages during synchronisation times, the available bandwidth of the network is temporarily but dramatically reduced. FTSP, in contrast, distributes the radio message load evenly and thus it provides constant radio bandwidth for applications.

7 Conclusions

Advances over the last few years in radio and wireless sensor network design enable, and increasingly challenging application requirements drive us, to continuously reconsider the boundaries between the hardware, operating system, middleware and application layers. Single layer solutions to several key problems that inherently touch multiple layers such as time synchronisation in wireless networks are likely to be less efficient than those that can carefully address the challenges at each layer. The price one must pay for the increased control and efficiency, however, is additional complexity between the layers. We argue that the solution is to minimise the number of different interfaces between the layers but increase their expressive power.

Our first key contribution is the careful analysis of the sources of time synchronisation jitter between the operating system, radio chip and wireless channel. We present a comprehensive list of software techniques and recommended hardware solutions to maximise the precision of the message time-stamping. Then we propose a time synchronisation primitive, called elapsed time on arrival (ETA) which is a one-way, sender-receiver time-stamping service for broadcasted messages. On the Mica2 platform its precision is better than 1.4 μs between the sender and/or any

of the receivers. It allows the translation of event times, expressed in the local clock of the nodes, from the sender to its neighbours. Other than the event time, it has absolutely no data overhead in the message. The two key enablers of ETA are 1) integration with the MAC-layer to minimise the jitter, and 2) the ability to embed the time-stamp in the message while it is being transmitted.

As our second contribution, we have identified a small set of canonical services together with their application programming interfaces that capture the time synchronisation requirements of actual wireless sensor network applications. We argue that abstracting common timesync usage patterns from the existing applications will help application developers to identify their time synchronisation needs. We propose to build canonical services on top of an explicit timesync primitive, such as our ETA, and show that their implementation with ETA is less complex and has comparable or better precision. The Event Time-Stamping canonical service is implemented by the Routing Integrated Time Synchronisation (RITS) protocol, which is the multi-hop extension of ETA. Combined with a convergecast routing policy, it allows the root node to correlate the detection time of a single event at multiple observers. In a 45-node 10-hop network the maximum and average time synchronisation errors were 80 μs and 8 μs , respectively. The Virtual Global Time canonical service was implemented by the Rapid Time Synchronisation (RATS) protocol, in which a root node broadcasts its local time in the network, using RITS. In a 60-node 10-hop network the algorithm achieved network-wide synchronisation in four seconds, and the maximum and average time synchronisation errors were 26 μs and 2.7 μs , respectively. In a similar setup, the Flooding Time Synchronisation Protocol (FTSP) achieved network-wide synchronisation in 10 minutes, while the maximum and average errors were 14 μs and 2.3 μs , respectively.

Acknowledgements

The DARPA/IXO NEST program has partially supported the research described in this paper.

References

- Blum, P., Meier, L. and Thiele, L. (2004) ‘Improved interval-based clock synchronization in sensor networks’, *Third International Symposium on Information Processing in Sensor Networks*, Berkeley, California, USA, pp.349–358.
- Dai, H. and Han, R. (1994) ‘Tsync: A lightweight bidirectional time synchronization service for wireless sensor networks’, *ACM SIGMOBILE Mobile Computing and Communications Review*, Vol. 8, No. 1, pp.125–139.
- Dutta, P.K. (2004) *On Random Event Detection in Wireless Sensor Networks*, Master’s Thesis, The Ohio State University, Columbus, Ohio, USA
- Elson, J. (2002) ‘Fine-grained network time synchronization using reference broadcasts’, *Proceedings of the Fifth Symposium on Operating Systems Design and Implementation (OSDI 2002)*, Boston, Massachusetts, USA, December.

- Elson, J. (2003) *Time Synchronization in Wireless Sensor Networks*, PhD Thesis, University of California, Los Angeles.
- Elson, J. and Romer, K. (2002) 'Wireless sensor networks: a new regime for time synchronization', *Proceedings of the First Workshop on Hot Topics in Networks (HotNets-I)*, Princeton, New Jersey, October 28–29.
- Ganeriwal, S., Kumar, R. and Srivastava, M.B. (2003) 'Timing-sync protocol for sensor networks', *ACM Sen-Sys 2003*, Los Angeles, California, USA.
- Hill, J. (2000) *A Software Architecture Supporting Networked Sensors*, Master's Thesis, U.C Berkeley Dept. of Electrical Engineering and Computer Sciences, San Francisco, California, USA.
- Hill, J. and Culler, D. (2002) 'Mica: a wireless platform for deeply embedded networks', *IEEE Micro*, Vol. 22, No. 6, pp.12–24.
- Horauer, M., Schossmaier, K., Schmid, U., Holler, R. and Kero, N. (2002) 'Evaluation of a high precision time synchronization prototype system for ethernet lans', *34th Annual Precise Time and Time Interval Meeting (PTTI)*, Reston, Virginia, USA, December.
- Romer, K., Blum, P. and Meier, L. (2005) 'Time synchronization and calibration in wireless sensor networks', in Stojmenovic, I. (Ed.): *Handbook of Sensor Networks: Algorithms and Architectures*, Wiley and Sons, pp.199–237.
- Kopetz, H. (1997) *Real-Time Systems: Design Principles for Distributed Embedded Applications*, Kluwer Academic Publishers, Norwell, Massachusetts, USA.
- Kopetz, H. and Ochsenreiter, W. (1987) 'Clock synchronization in distributed real-time systems', *IEEE Transactions on Computers*, Vol. C-36, No. 8, pp.933–939.
- Kopetz, H. and Schwabl, W. (1989) *Global Time in Distributed Real-time Systems*, Technical Report 15/89, Technische Universitat Wien, Austria.
- Lamport, L. (1978) 'Time, clocks, and the ordering of events in a distributed system', *Commun. ACM*, Vol. 21, No. 7, pp.558–565.
- Ledeczi, A., Volgyesi, P., Maroti, M., Simon, G., Balogh, G., Nadas, A., Kusy, B., Dora, S. and Pap, G. (2005) 'Multiple simultaneous acoustic source localization in urban terrain', *Proc. 4th International Symposium on Information Processing in Sensor Networks (IPSN 2005)*, Los Angeles, California, USA, April.
- Li, Q. and Rus, D. (2004) 'Global clock synchronization in sensor networks', *IEEE INFOCOM*, Hong Kong, China, March.
- Maroti, M. (2004) 'Directed flood-routing framework for wireless sensor networks', *Proc. 5th ACM International Middleware Conference*, Toronto, Ontario, Canada, pp.99–114.
- Maroti, M., Kusy, B., Simon, G. and Ledeczi, A. (2004) 'The flooding time synchronization protocol', *Proc. 2nd ACM International Conference on Embedded Networked Sensor Systems (SenSys 2004)*, Baltimore, Maryland, USA, pp.39–49.
- Marzullo, K. and Owicki, S. (1983) 'Maintaining the time in a distributed system', *Second Annual ACM Symposium on Principles of Distributed Computing*, pp.295–305.
- Meier, L., Blum, P. and Thiele, L. (2004) 'Internal synchronization of drift-constraint clocks in ad-hoc sensor networks', *Fifth ACM International Symposium on Mobile Ad Hoc Networking and Computing*, Tokyo, Japan, pp.90–97.
- Reichenbach, H. (1957) *The Philosophy of Space and Time*, Dover, New York.
- Romer, K. (2001) 'Time synchronization in ad hoc networks', *Proceedings of the International Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc '01)*, Long Beach, California, USA, October 4–5.
- Schmid, U. and Schossmaier, K. (1997) 'Interval-based clock synchronization', *Real-time Systems*, Vol. 12, No. 2, pp.172–228.
- Sichitiu, M.L. and Veerarittiphan, C. (2003) 'Simple, accurate time synchronization for wireless sensor networks', *IEEE Wireless Communications and Networking Conference (WCNC03)*, New Orleans, Louisiana, USA, March 16–20.
- Sivrikaya, F. and Yener, B. (2004) 'Time synchronization in sensor networks: A survey', *IEEE Network*, Vol. 18, No. 4, pp.45–50.
- Su, W. and Akyildiz, I.F. (2004) 'Time-diffusion synchronization protocol for sensor networks', To appear in *IEEE/ACM Transactions on Networking*, Vol. 13, No. 1.
- van Greunen, J. and Rabaey, J. (2003) 'Lightweight time synchronization for sensor networks', *2nd ACM International Workshop on Wireless Sensor Networks and Applications*, pp.11–19.
- Xu, N., Rangwala, S., Chintalapudi, K., Ganesan, D., Broad, A., Govindan, R. and Estrin, D. (2004) 'A wireless sensor network for structural monitoring', *Proc. 2nd ACM International Conference on Embedded Networked Sensors Systems (SenSys 2004)*, Baltimore, Maryland, USA, pp.13–14.
- Chipcon (2003) *Chipcon CC2420 Datasheet: 2.4 GHz IEEE 802.15.4/ZigBee-ready RF Transceiver (v 1.2)*, Chipcon SA, Oslo, Norway.