

Lecture on Sensor Networks

Copyright (c) 2005 Dr. Thomas Haenselmann (University of Mannheim, Germany).
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.2
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts.
A copy of the license is included on the last pages entitled "GNU
Free Documentation License".

Synchronization in sensor networks

Motivation

Systeme und Metriken zur Zeitmessung

TAI: Temps Atomique International ist eine gemittelte Zeit aus mehreren Cäsium Atomuhren, die über die ganze Welt verteilt sind. Das Cäsium wechselt seinen elektrischen Zustand 9.192.631.770 mal/Sekunde. Diese Zustandswechsel werden gezählt, um den Ablauf einer Sekunde zu messen. Dabei entsteht eine Abweichung von etwa 10^{-12} pro Uhr, die gemittelte Zeit mehrerer Atomuhren ist entsprechend genauer.

UTC: Temps Universel Coordonné ist die weltweit vereinbarte korrigierte Zeit, die sich an die TAI anlehnt, jedoch in manchen Jahren um eine Schaltsekunde korrigiert werden muss, da die Erdumdrehung durch die Gezeiten abgebremst wird, sich der Mond von der Erde entfernt, die Sonne kleiner wird etc.

NTP: Network Time Protocol mit dem NTP Zeitstempel, der aus 64 Bit besteht. Dabei stehen die ersten 32 Bit für einen Sekundenzähler, ab dem 1. Januar 0:00 Uhr im Jahr 1900, die letzten 32 Bit bezeichnen den Bruchteil einer Sekunde, d. h. mit einer Genauigkeit von $1/2^{32}$ Sekunden.

Time and error types

Reference Broadcast Synchron.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Motivation

Naive Lösung für Synchronisation

Anfrage bei einem Knoten, der selbst die genaue Zeit hat und Korrigieren der eigenen Uhr. Ist die zu erwartende Laufzeit der Nachricht klein, vor allem im Verhältnis zur benötigten Genauigkeit der Uhr, so reicht die einfache Nachfrage aus. Sonst kann man die Latenz zwischen Sender der Uhrzeit und Empfänger wie folgt unterteilen:

- Sendezeit:** Die Zeit, die der Sender braucht um die Nachricht zu erstellen, ein Multitasking eine Zeitscheibe zu bekommen, evtl. bestehende Warteschlangen abzubauen etc.
- Zugriffszeit:** Gemeint ist die Zeit die vergeht, bis Zugriff auf das Medium gewährt werden kann (siehe CSMA Verfahren etc.) Insbesondere in Sensornetzen bei allen Verfahren mit Contention-Phase wie S-MAC. T-MAC etc.
- Ausbreitungszeit:** Die Zeit, die die Nachricht selbst auf dem Medium braucht. Zwischen Sendern vernachlässigbar, über das Netzwerk hinweg können sich jed. lange Latenzzeiten akkumulieren, vor allem bei großer Anzahl Hops.
- Empfangszeit:** Analog zur Zeit für das Senden, d. h. gemeint ist die Latanz, bis die Nachricht in der Anwendungsschicht angelangt ist.

Time and error types

Reference Broadcast Synchron.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation


Comparison of time intervals

Synchronization in sensor networks

Motivation

Reference broadcast synchronization (RBS)

Principle: Everyone gets the same synch signal over the broadcast channel. Despite the different signal propagation times for nodes with varying distances to the synchronizer every node is triggered at the same time. The synch signal does not have to be meaningful. A pulse or any kind of even can be used.



The signal propagation time is (almost) negligible. The source for variance in the measurement is the time different nodes need to receive the signal. All kinds of delays at the sender side do not play an important role, since they are the same for every receiver.

Time and
error types

Reference Broad-
cast Synch.

Round Trip Time
based sync.

Error in drift
estimation

Time sync
Protocol

Interval
estimation

Comparison of
time intervals

Synchronization in sensor networks

Motivation

Reference broadcast synchronization (RBS)

When a sync signal is received every receiver stores its own local time. No synchronization took place so far. However, when exchanging timestamps with peer nodes next time the drift between clocks can be estimated.

Example: Node A receives a broadcast signal at 17:05. Later B tells A about an event that happened at 17:35 and adds in addition that the last sync signal was received at 17:02 according to B's clock. So A can conclude that its own clock is 3 minutes ahead of B and that the event took place at 17:40 according to its local time.

In addition B could adjust its clock to be compliant with A if A's time is more credible.

Advantage: The broadcast signal can be any event which can be received by a group of nodes, e.g., even an RTS or CTS. But it has to be clear that everyone is talking about the same event when calculating the drift. An increasing counter for the synchronization can be helpful in this context.

Disadvantage: All receivers have to be in the same domain (listening to the same channel).

Time and error types

Reference Broadcast Synch.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

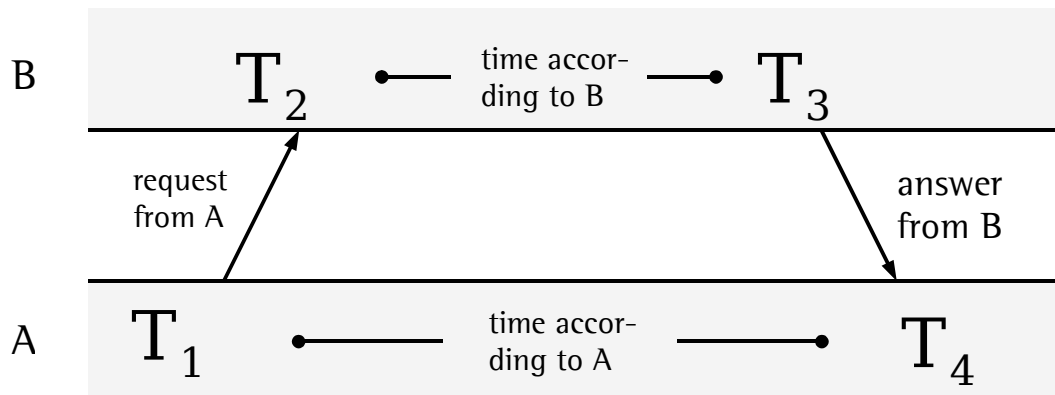
The Network Time Protocol (NTP) - Round Trip Time (RTT) based

Estimation of the latency delta (packet delay between to nodes) and the clock drift theta.

$$\delta = \frac{(T_4 - T_1) - (T_3 - T_2)}{2}$$

$$\theta = \frac{(T_2 - T_1) + (T_3 - T_4)}{2}$$

theta > 0 means that clock B is ahead of clock A



Time and error types

Reference Broadcast Sync.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

The Network Time Protocol (NTP)

Estimation of the round trip time delta and clock drift theta

$$\delta = \frac{(T_4 - T_1) - (T_3 - T_2)}{2} \quad \theta = \frac{(T_2 - T_1) + (T_3 - T_4)}{2}$$

Verification: The clock shall deviate about e time units and the single packet delay (1 hop communication) is assumed to be d time units.

$$T_2 = T_1 + d + e$$

$$T_4 = T_3 + d - e$$

$$\delta = \frac{(T_3 + d - e - T_1) - (T_3 - (T_1 + d + e))}{2} = d$$

$$\theta = \frac{(T_1 + d + e - T_1) + (T_3 - (T_3 + d - e))}{2} = e$$

Time and error types

Reference Broadcast Synchron.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Error analysis of clock drift estimation between RTT and RBS

RTT based approach: The round trip time (time to send a packet and get an answer) is used to estimate the true delay of a simple one-way communication and the clock drift.

RBS based approach: A number of receivers are synchronized by the same event which is used to estimate the clock drift. Once the drift is known the packet delay can also be derived.

D_K = Drift between K's clock and real time

S_K = time to **Send** (to prepare the packet at the sender, wait for empty queue etc. but not the time to access the medium).

A_K = Waiting time to gain **Access** to the medium (an issue of the MAC protocol)(Access)

$P_{K,L}$ = Signal **Propagation** time between node K and L

R_L = Time to **Receive** a packet (basically to hand it though the network stack to the application layer.

T_1 = This is the time as it is perceived by node 1.

t_1 = The actual realtime belonging to T_1 . So the drift is obtained by $T_1 - t_1$.

D_1 = Drift of node 1.

The length of the message itself and the time to send the whole message is not an issue of the above estimation. We only consider the arrival of the first bit.

Time and error types

Reference Broadcast Synch.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Error analysis of clock drift estimation between RTT and RBS

Error in drift estimation of RTT based approaches:

$$T_1 = t_1 + D_A$$

$$T_2 = T_1 + S_A + A_A + P_{A,B} + R_B$$

$$T_3 = t_3 + D_B$$

$$T_4 = T_3 + S_B + A_B + P_{B,A} + R_A$$

Drift:
$$\theta_1 = \frac{(T_2 - T_1) + (T_3 - T_4)}{2}$$

Error in drift estimation: We substitute the measured times T with the actual and true real time t, however extended by all errors which can occur in the communication. In the end we want to know which kinds of errors influence the drift calculation most and which drop out.

$$e_{\theta_1} = \frac{(t_1 + D_A + S_A + A_A + P_{A,B} + R_B - (t_1 + D_A)) + (t_3 + D_B - (t_3 + D_B + S_B + A_B + P_{B,A} + R_A))}{2}$$

$$e_{\theta_1} = \frac{S_A + A_A + P_{A,B} + R_B - S_B - A_B - P_{B,A} - R_A}{2}$$

We can assume that the propagation $P_{A,B}$ from node A to B is approximately equal to the inverse path $P_{B,A}$ so that the propagation time compensate one another.

$$e_{\theta_1} = \frac{S_A + A_A + R_B - S_B - A_B - R_A}{2}$$

When using nodes of the same kind the expectation values of the errors should have the same degree of magnitude. So the degree to which they cancel out one another depends on their variance. A smaller variance means better error cancellation.

Time and error types

Reference Broadcast Sync.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Error analysis of clock drift estimation between RTT and RBS

Error in drift estimation of RBS based approaches:

$$T_1 = t_1 + D_S$$

$$T_2 = T_1 + S_S + A_S + P_{S,A} + R_A$$

$$T_3 = T_1 + S_S + A_S + P_{S,B} + R_B$$

Drift:

$$\theta_2 = T_2 - T_3$$

Node S sends a synch signal at time T_1 (the time S measured with its clock). A receives the signal at T_2 , B at T_3 (according to their local clock respectively).

Error in drift estimation:

$$e_{\theta_2} = T_2 - T_3 = t_1 + D_S + S_S + A_S + P_{S,A} + R_A - (t_1 + D_S + S_S + A_S + P_{S,B} + R_B)$$

$$e_{\theta_2} = P_{S,A} + R_A - P_{S,B} - R_B$$

As compared to the error in the RTT based approach:

$$e_{\theta_1} = \frac{S_A + A_A + R_B - S_B - A_B - R_A}{2}$$

All latencies at the sender side cancel off themselves at the receiver. In contrast the signal propagation time is not equal here anymore since the receivers may have different distances to the synchronizer. Furthermore the term is not divided by 2. On the other hand many erroneous influences are missing and the signal propagation time is very small.

Time and error types

Reference Broadcast Synch.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

TPSN: Timing-sync Protocol for Sensor Networks

(according to the publication of S. Ganeriwal, R. Kumar, M. B. Srivastava)


Initialization phase

The initialization originates in TPSN from the highest level, which is the root node in the beginning. It sends a `level_discovery` packet including its own level.

All nodes within the range receive this announcement and adopt the level after increasing it by one. Then every node starts its own random timer.

After a timer's expiration a node will also send a `level_discovery` packet which is directed to all nodes in the range that did not yet hear an announcement. The others will ignore it. At the end of the flooding process every node has obtained a level which is increasing with an increasing number of hops to the root node.

Note that it is not guaranteed that a node has obtained its level because the `level_discovery` messages of more than one sender could have collided without the senders knowing it (remember the hidden station problem). A node with no valid level can send a `level_request` packet which is answered by everyone after the expiration of a random timer on the senders' sides. The requesting node will then choose the smallest level and increase it by 1.



Time and error types

Reference Broadcast Sync.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Motivation

Time-Sync Protocol for sensor Networks (TPSN)


Synchronization phase

The root node issues a `time_sync` packet which starts a random timer at the level 1 nodes. On expiration the node with the shortest random timer asks the root for synchronization using a **synchronization pulse**. Synchronization always takes place between a child and a parent node.

After the synchronization pulse arrived at the parent node it will reply with an **ACK**. Now the requesting node knows the round trip time and thus can estimate the single packet delay in the 1-hop communication.

The nodes on a higher level overhear the **synchronization pulse** however they start a random timer themselves. It has to be chosen large enough so that the synchronization between the parent and grand parent node can be finished. On expiration of the first timer on the second level the resp. node will send a **synchronization pulse** to its father node and so on. Sync pulse and ACK have the same role as ping and pong in NTP.

The synch pulse has the side effect that a node can recognize that it did not receive one and that the child nodes may not be synchronized. In this case a parent node will simply send a `time_sync` packet itself (this is usually only done by the root) to trigger its children. The `time_sync` packet is however not necessary in the regular operation.



Time and error types

Reference Broadcast Sync.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Estimation of time intervals in sensor networks

Precise synchronization and consistency of clocks are not always realistic in sensor networks due to the increased overhead or

- no synch source available
- clocks not precise due to large temperature deviations.
- timing errors accumulate over several clocks
- clocks may not only experience a drift but also a so-called slew (clock too fast about a factor)

An event is detected in the time domain of the observer. Rather than trying to adapt it to a global time the inaccuracy is accepted. When telling another node N the event is simply transformed into the time domain of N. So the knowledge propagates through the local time domain of a number of hops.

Example: An observer detects an event at 19:00 and tells it to B. B's clock shows 18:00. So B knows about the same event at the approximately same time but calls that time 18:00 rather than 19:00. Then the message is even propagated further to C which calls the as time 19:00 again according to its clock. The fact the B called it 18:00 in between is not a problem.

Problem: When forwarding the message between hops some additional time passes. These latencies have to be considered in the above example. One of the most important latencies is the packet delay which is often estimated as $0.5 \times$ round trip time. Note that RTT means the way to another node and the way back. We do not know how much time was actually spend on a single hop. The lower bound is zero (unrealistic? -> find a better lower bound) or the whole RTT for one hop as upper bound.

The remaining uncertainty can be accounted for by defining an interval with a lower bound for the event and an upper bound.

Time and error types

Reference Broadcast Sync.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Estimation of time intervals in sensor networks

Approach for estimating time intervals for events according to „Time Synchronization in Ad Hoc Networks“ by Kay Römer

Node 1:

$$[r_1, r_1] = [T_{\text{Messung}}, T_{\text{Messung}}]$$

Node 2:

$$\left[r_2 - (s_1 - r_1) \frac{1+p_2}{1-p_1} - \left(rtt_2 - idle_1 \frac{1-p_2}{1+p_1} \right), r_2 - (s_1 - r_1) \frac{1-p_2}{1+p_1} \right]$$

Node 3:

$$\left[r_3 - (s_1 - r_1) \frac{1+p_3}{1-p_1} - (s_2 - r_2) \frac{1+p_3}{1-p_2} - \left(rtt_2 \frac{1+p_3}{1-p_2} - idle_1 \frac{1-p_3}{1+p_1} \right) - \left(rtt_3 - idle_2 \frac{1-p_3}{1+p_2} \right), \dots \right. \\ \left. \dots r_3 - (s_1 - r_1) \frac{1-p_3}{1+p_1} - (s_2 - r_2) \frac{1-p_3}{1+p_2} \right]$$

Node N:

$$\left[r_N - (1+p_N) \sum_{i=1}^{N-1} \frac{s_i - r_i}{1-p_i} - (1+p_N) \sum_{i=1}^{N-1} \frac{rtt_i}{1-p_i} + (1-p_N) \sum_{i=2}^N \frac{idle_i}{1+p_{i-1}} + rtt_N, r_N - (1-p_N) \sum_{i=1}^{N-1} \frac{s_i - r_i}{1+p_i} \right]$$

Exceptional: rtt_3 does not have to be transformed. Because it was measured by node 3 itself so it already is in the time domain of node 3.

Time and error types

Reference Broadcast Synch.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

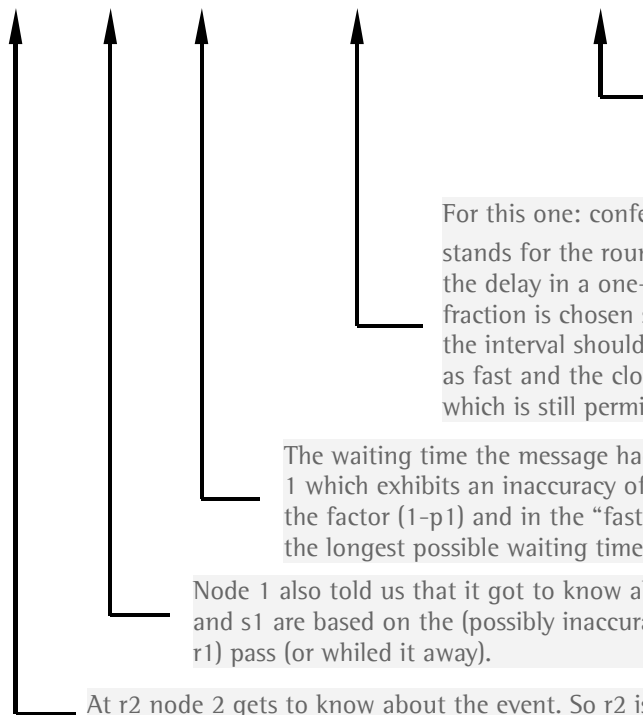
Comparison of time intervals

Synchronization in sensor networks

Estimation of time intervals in sensor networks

Node 2:

$$\left[r_2 - (s_1 - r_1) \frac{1 + p_2}{1 - p_1} - \left(rtt_2 - idle_1 \frac{1 - p_2}{1 + p_1} \right), r_2 - (s_1 - r_1) \frac{1 - p_2}{1 + p_1} \right]$$



In contrast to the lower bound the upper should remain as large (=as young) as possible. We do not want to subtract more than necessary. So we chose the lower estimate for the packet delay which is 0.*1. The message's waiting time in node 1 should also be as small as possible which means a fast clock at node 1 and a slow one at node 2.

For this one: confer the explanation on the next but one slide. $(rtt_2 - idle_1 \times \text{fraction})$ stands for the round trip time between node 1 and 2 and exhibits an upper bound for the delay in a one-hop communication (the expectation value however, is $RTT/2$). The fraction is chosen such as to yield a large overall latency. Because the lower bound of the interval should become as small as possible. Therefore the clock of node 1 has to be as fast and the clock of node 2 has to be as slow as possible – to the maximum degree which is still permitted by the clock's known inaccuracies.

The waiting time the message had to pass at node 1 is based on the clock (or time domain) of node 1 which exhibits an inaccuracy of $\pm p_1$. So in the "slowest possible case" the clock is too slow by the factor $(1-p_1)$ and in the "fastest" case it is too fast by the factor $(1+p_1)$. The quotient results in the longest possible waiting time at node 1 – based on the time domain of node 2.

Node 1 also told us that it got to know about the event at r_1 and that it informed us at s_1 . Note that both r_1 and s_1 are based on the (possibly inaccurate) clock of node 1. We could say that node 1 let the time span $(s_1 - r_1)$ pass (or whiled it away).

At r_2 node 2 gets to know about the event. So r_2 is a first estimate for the time of the event. In any case it can not have happened later than r_2 .

*1 Why can the packet delay simply be omitted? Because we only know the round trip time. However, we do not know how this overall RTT splits into the way to another node and the way back. The only obvious lower bound is that one way did take no time (even though this is not realistic) and the only obvious upper bound is that the packet took the full RTT for one way (leaving no time for the way back. Any better ideas?).

Time and error types

Reference Broadcast Sync.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Estimation of time intervals in sensor networks

What's the matter with the calculation (rtt-idle)?

The receiver E wants to estimate an upper bound for the one-way packet delay between itself and the sender. Therefore, it could simply send the NTP typical ping and pong message. But this would require two additional packets which is not a good option in terms of energy efficiency.

Solution: Often E has sent S a message before for other reasons at time T_1 . This message can in a sense already be interpreted as an NTP like ping message. Both E and S stored the local time when the message was sent (T_1) resp. when it was received (T_2).

Later at time T_3 node S sends a message to E. Prior to preparing the message it looks up in its neighborhood table when it received the last message from E, which is the time T_2 in our case. So it adds to the message that the **idle period** between itself and node E lasted for $T_3 - T_2$ time units. This message is then received by E at time T_4 . E also looks up in its table that it did communicate with S at last time at T_1 . So the round trip time is $(T_4 - T_1) - (T_3 - T_2)$. The process does not differ much from the NTP ping and pong messages.

Disadvantage: S and E have to do some simple book keeping for every neighbor they have. Furthermore a long idle period can introduce an additional error which can even be in the order of magnitude of the RTT itself for large idle times.

Time and error types

Reference Broadcast Synch.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Estimation of time intervals in sensor networks

Comparison of time intervals


Points in time can be compared easily, e.g., to evaluate whether one event occurred before another. For intervals the comparison becomes slightly more complicated.

We assume that E_1 lies in the interval $[t_1, t_2]$, E_2 in the interval $[t_3, t_4]$.

Did E_1 happen before E_2 ?

yes	if $t_2 < t_3$
no	if $t_1 > t_4$
maybe	else

Closer analysis of maybe: What is the probability of E_1 happening before E_2 ?



We can distinguish three cases. The required order can be ensured if either $E_1 < t_3$ or $E_2 > t_2$. If both happen in the common time span $[t_3, t_2]$ there is a 50% chance of the required order.

Time and error types

Reference Broadcast Synch.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

Synchronization in sensor networks

Estimation of time intervals in sensor networks

Comparison of time intervals

But it is much easier to evaluate the inverse question: Is $E_1 > E_2$? This is the case if $E_1 > t_3$...

$$P(E_1 > t_3) = \frac{t_2 - t_3}{t_2 - t_1}$$

... and $E_2 < t_2$:

$$P(E_2 < t_2) = \frac{t_2 - t_3}{t_4 - t_3}$$

Since both events have to happen at the same time the product of the two probabilities have to be computed. Finally, it E_1 and E_2 are in the same interval there can be two different orders, but only one is valid in this context. Dividing by 2 takes this into account.

$$P(E_1 > E_2) = \frac{(t_2 - t_3)(t_2 - t_3)}{2(t_2 - t_1)(t_4 - t_3)}$$

In the beginning we aimed at obtaining the probability of $E_1 < E_2$.

$$P(E_1 < E_2) = 1 - P(E_1 > E_2)$$

Time and
error types

Reference Broad-
cast Synch.

Round Trip Time
based sync.

Error in drift
estimation

Time sync
Protocol

Interval
estimation

**Comparison of
time intervals**

Synchronization in sensor networks


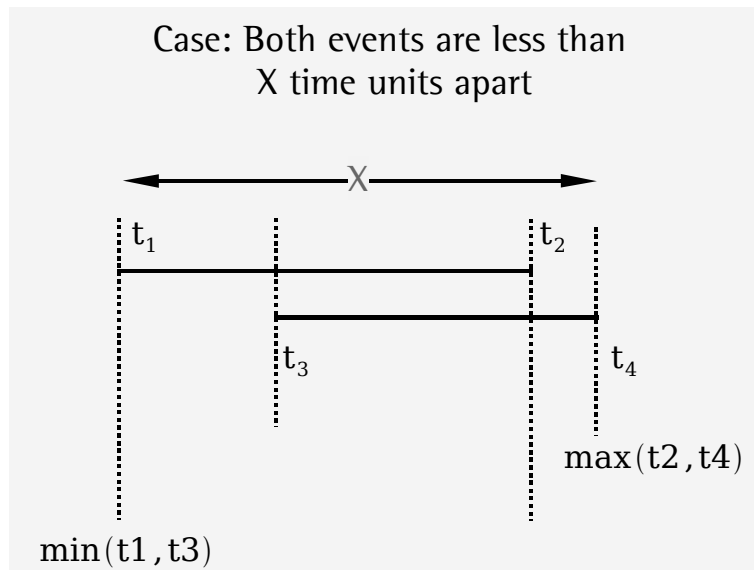
Estimation of time intervals in sensor networks

Comparison of time intervals

Are E_1 and E_2 further apart than X time units?

yes	if $\max(t_4, t_2) - \min(t_3, t_1) < X(1-p)$
no	if $\max(t_3, t_1) - \min(t_2, t_4) > X(1+p)$
maybe	else

If it is known in advance that the clock's deviation come from the interval $[1-p, 1+p]$ the real time X can be adjusted according to the question.



Time and error types

Reference Broadcast Sync.

Round Trip Time based sync.

Error in drift estimation

Time sync Protocol

Interval estimation

Comparison of time intervals

GNU Free Documentation License
Version 1.2, November 2002

Copyright (C) 2000,2001,2002 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties; any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

son of
intervals

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.
- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.