Same time – different place

Precise time synchronization for automation purposes

Tor Skeie, Svein Johannessen, Trond Løkstad, Øyvind Holmeide

For most people, synchronizing their watch to within one second of the radio's time signal is close enough. Not in substation automation, however. Here, the required accuracy is more likely to be of the order of one microsecond – precision that calls for dedicated, expensive equipment.

To solve its synchronization problems industry is turning to the switched, fast Ethernet. There is just one critical question: Can this communications standard guarantee the required precision? The answer is yes. ABB and OnTime Networks AS have joined forces to develop solutions that achieve IEC class T5 (1 µs) and T3 (25 µs) time synchronization over switched Ethernet.

Recently, the concept was implemented in AC 800M controllers for the prominent Grane oilfield project, where it provides high-precision synchronization in ABB's Industrial^{IT} architecture.

Time stamping is ubiquitous in the PC world. Every document saved is automatically assigned a date and time value. This simplifies listing as well as retrieval, but in each case the accuracy required is of the 'wristwatch' order – a second or so.

If the PC is now connected to a network and starts exchanging documents and e-mails, time stamping becomes a little more demanding. Wristwatch standard is still enough, but the PCs must be *synchronized*. Otherwise, your reply to an e-mail could carry a time stamp that precedes the original e-mail's.

A solution is to elect one PC as the 'timekeeper' from which the other PCs get their time. This solution works well on a local office setup – the variation in time lag from the timekeeper to an individual PC is relatively small – but not for Internet applications.

Although its first use is in an offshore application, the new synchronization concept has benefits for many industrial sectors.

Enter the Internet Network Time Protocol. With this a time-stamped time request message is sent to a 'timeserver', which adds an arrival time stamp and a retransmit time stamp before returning the request message to the requesting PC. The requesting PC time stamps the message when it returns and uses all the time stamps in calculating the correct time. This protocol and its little brother, the Simple Network Time Protocol, are able to synchronize computers across the Internet with a precision in the low milliseconds.

Automation and the need for networking

Traditional automation and measurement systems make use of the 'fieldbus' concept for their networking tasks. This concept, however, suffers from a fragmented market, resulting in low volumes and, therefore, high prices. Using switched, fast Ethernet for automation purposes therefore looked like being a sure winner – providing some practical problems could be solved:

- Ethernet is by nature 'democratic' – all nodes have an equal chance of accessing the network at a given point in time. Automation systems are 'dictatorial' – some nodes are more important than others and want to be given priority when accessing the network.
- Ethernet is by nature 'statistical' – there is no way of predicting just when a queued message will actually arrive at the destination. Automation systems want assurance, not probabilities.
- Ethernet is a full-fledged communication network with a fairly large overhead – small data packets (like the ones in traditional data collection) are grossly inefficient.

These problems are being addressed [1]. Switched Ethernet goes a long way toward guaranteeing a maximum message delivery time [2]. Intelligent nodes with local real-time clocks sample their data at predefined points in time, time stamp

1 Future substation communications

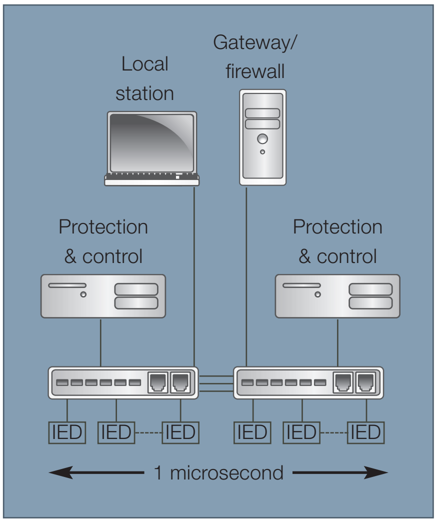

Communication network in a future substation automation system, with Ethernet used as a common network infrastructure.

the data, pack several data sets into a network packet, and transmit the data at their leisure. The main thing missing is the connection between the local clock and the one in the central controller. Before giving this more thought, it is worth having a look at the automation field with the greatest need for accurate time synchronization – substation automation.

Synchronization requirements in process automation

Substation automation (SA) demands very strict synchronization accuracy, with different protection functions requiring different levels of precision (see table).

The substation automation sector is currently about to migrate to switched, fast Ethernet for the automation network infrastructure. The ultimate vision here is full interoperability between products from different vendors at all levels. The introduction of *de-facto* standard concepts and adoption of off-the-shelf technologies are the key instruments that will enable this goal to be reached 1.

Several studies have shown that switched, fast Ethernet has sufficient real-time characteristics to meet SA demands [3]. What is left to show is that it is possible to implement the various IEC classes of synchronization accuracy. This is considered to be the final obstacle to full migration to Ethernet in SA. Currently, to ABB's knowledge, no one has demonstrated the possibility of highly accurate time synchronization over switched Ethernet. Such a step would significantly reduce the cabling and transceiver cost, since dedicated links are used for this purpose today.

Time synchronization status

Network Time Protocol (NTP) is the most prominent public domain synchronization method [4]. It relies on sophisticated mechanisms to access national time to organize time synchronization subnets, possibly implemented over various media, as well as adjustment of the local clock in each participating peer. SNTP targets simpler synchronization purposes.

These are both used for WAN and LAN synchronization where millisecond accuracy suffices. However, they are not solutions for accuracy in the low microsecond range.

A look at the automation field in general, and the SA world in particular, shows that there is a diversity of proprietary and patented solutions (including some from ABB) that achieve this accuracy. Nevertheless, an open, transparent and globally available and accepted solution would be better.

Why network synchronization is difficult

The delay between time stamping of a message in the source node and its time stamping in the destination node can have many causes: message preparation and handling, communication stack traversal (transmission and reception) and network access, among them. Variations in these individual delays can be due to unpredictability in the real-time operating system (RTOS) scheduling and network access and variations in the network transversal time. Time stamping at the lowest stack level helps to eliminate the stack delay variations and RTOS scheduling unpredictability, but it also introduces some complications in the implementation.

Extreme accuracy with IEC class T3

As mentioned above, much of the relatively low precision of NTP/SNTP implementations on a LAN stems from the time stamping of incoming and outgoing NTP/SNTP packets at the NTP/SNTP application layer. This makes time stamping a victim of RTOS scheduling granularity and unpredictability.

A look will now be taken at how 25 μs accuracy can be achieved using a tuned SNTP implementation and standard Ethernet switches.

A tuned SNTP time protocol implementation

The fact that an RTOS guarantees a maximum response time for an event but allows a wide variation below that time introduces a substantial variation in the time spent in the communication stacks. This has necessitated interrupt-level time stamping both in the time client and timeserver. The IEC class T3 solution described here adheres to the principle of interrupt-level time stamping of the SNTP request packet when sent from the time client and when received at the timeserver.

Moreover, ABB proposes that the synchronization be based on the transmit time stamp set by the client (referred to as T1 in SNTP terminology) and the receive time stamp set by the server (T2). Use of possible low-level transmit time stamping of the corresponding SNTP reply packet (T3) necessitates novel techniques for controlling the non-deterministic transmission of an Ethernet packet on the Ethernet.

EC class	Synchronization accuracy
T1	1 ms
T2	0.1 ms
T3	±25 μs
T4	±4 μs
T5	±1 μs

2 Client and server

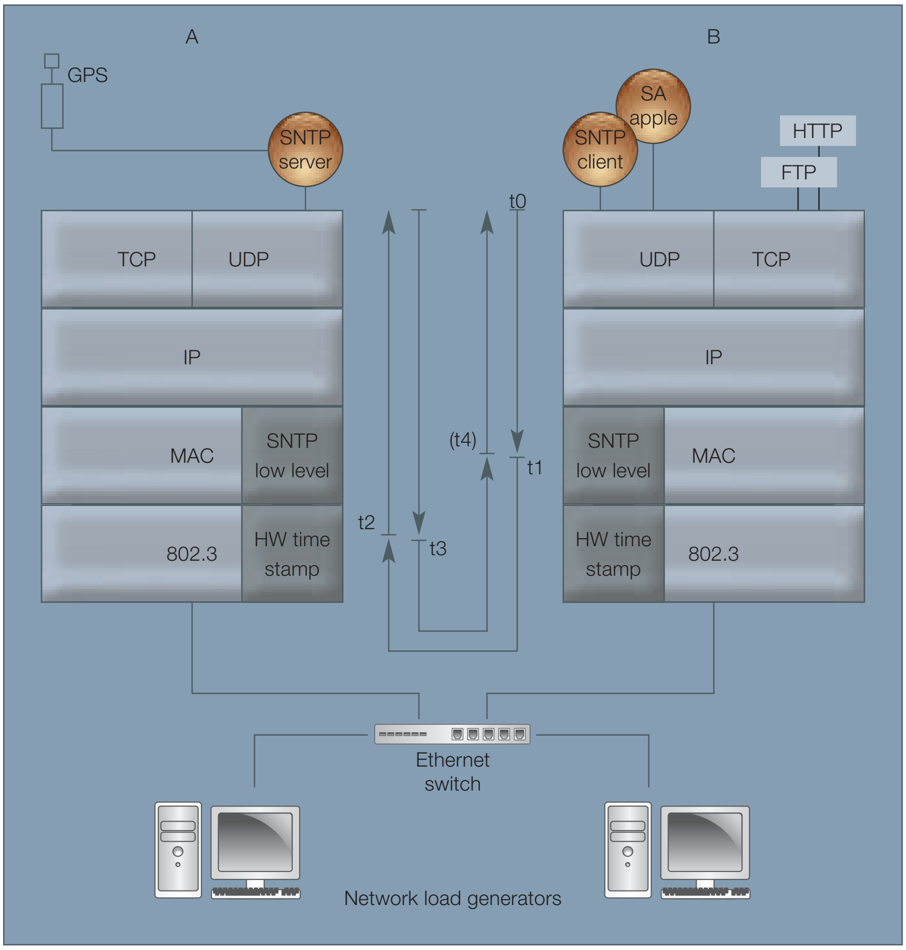

SNTP time client (B) – server (A) relationship with low-level transit time stamping.
MAC = Media Access Control UDP = User Datagram Protocol

A side-effect of using just T1—T2 is that no mechanism for automatic calibration of the network latency is then available, making it necessary to manually calibrate the propagation delays of the drop links and the minimum switch latency **2**.

Time client implementation issues

In the time client three different low-level time stamping methods are evaluated:

- Hardware time stamping in the Ethernet controller
- Software time stamping in an Interrupt Service Routine (ISR) outside the RTOS. This ISR should be connected to the Ethernet Interrupt Request signal and have a top hardware priority.
- Software time stamping in an Interrupt Service Routine (ISR) controlled by the RTOS (Ethernet driver). This ISR is connected to the Ethernet Interrupt Request signal with a normal hardware priority.

Only the first two of these prove to be suitable for very accurate time synchronization.

> Hardware time stamping or low-level software time stamping outside the RTOS eliminates the client inaccuracy from the error budget of the SNTP time synchronization loop.

Use of any of the above low-level time stamping methods is considered an implementation issue and will not cause any incompatibly between a low-level time stamping client and a standard high-level time stamping server. In addition to low-level time stamping, the time client must consider the following aspects:

- The interval between time updates.
- The specifications of the local time-of-day clock with respect to resolution, accuracy/stability, and the availability of drift and offset correction mechanisms.
- Use of adaptive filtering and time stamp validation methods in order to remove network delay variations.

Timeserver implementation issues

The timeserver should preferably implement one of the first two time stamping techniques listed above. At the very least it should be able to time stamp an incoming message with an accuracy of better than 2 µs, irrespective of the network load. The timeserver that was tested, from OnTime Networks 0, used hardware time stamping. The exact time should be taken from a GPS receiver, and the time parameters distributed from the timeserver should be based on GPS time representation instead of absolute time (ie, UTC timing) in order to cope with the 'leap second' problem. The timeserver should also support full duplex connectivity in order to avoid a situation where upstream data introduces extra switch latency in downstream data (ie, time requests).

Ethernet infrastructure implementation ssue

Preferably, there should be only one switch between a time client and a timeserver. Multiple switch levels will increase jitter (variations in the delay) throughout the infrastructure, which again might call for more complex filtering on the time client side. The Ethernet switch must also have good switch latency characteristics. The switch latency from the client drop link to the server drop link depends on several parameters:

- General switch load: This refers to all network load on the switch except for the packets sent to the timeserver. The variations in switch latency from the client drop link to the server drop link should be less than 2 µs.
- Timeserver load: This refers to other packets sent to the timeserver that may introduce extra delay in the transmission of a given SNTP request packet.
- Store-and-forward or cut-through: Most switches are based on 'store-and-forward' technology in which the whole Ethernet packet is received (and CRC checked) before it is passed through the switch. The delay then depends on the length of the packet. This is a static parameter for a given switch, specified by the vendor.

Tests on an actual network

Measurements were carried out in two steps. First, several Ethernet switches from different vendors were evaluated with respect to latency variations. Then, measurement data from a full network test setup verifying the various implementation options were collected. The conclusions drawn from the switch experiments were:

- Traffic not destined for the timeserver does not interfere with traffic to the timeserver.
- The switch latency for Ethernet packets to the timeserver depends to a great extent on other traffic to the timeserver.

Based on the results of the full network measurements, it was concluded that:

- Software time stamping using a sufficiently high priority interrupt (preferably non-maskable) is for all practical purposes indistinguishable from time stamping with special-purpose hardware.
- Software time stamping using an interrupt under RTOS control needs sophisticated filtering and statistical techniques before it can be used for time synchronization purposes. Because of this, such a method must be considered unsuitable for IEC class T3 synchronization.
- IEC Class T3 time synchronization using tuned SNTP over a switched Ethernet has been shown to be perfectly feasible.

3 Integrated timeserver

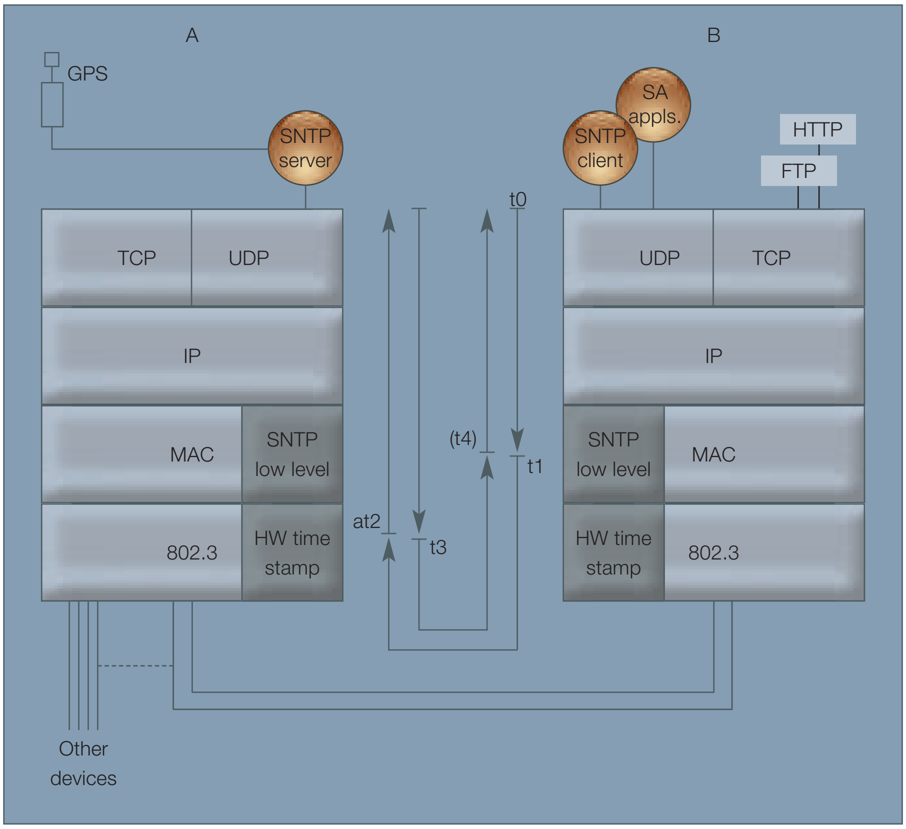

ABB PISA with SNTP time client (B) makes use of an Ethernet switch (A) with integrated timeserver.
PISA = Process Interface for Sensors and Actuators

Beyond the speed of light – IEC class T5

OnTime Networks [6] provide industrial-class fast Ethernet switches which are eminently suitable for substation automation. They are intended for applications with real-time requirements and therefore offer some additional features:

- Real-time data may be transmitted using packet priority (for example IEEE 802.1p) in order to guarantee worst-case switch latency.
- Each switch may be delivered with an integrated SNTP timeserver.

> OnTime solved the non-deterministic Ethernet bus access problems for SNTP request packets by integrating the SNTP timeserver into industrial-class fast Ethernet switches

Features provided by the SNTP timeserver in OnTime switches are:

4 The Grane oilfield installation

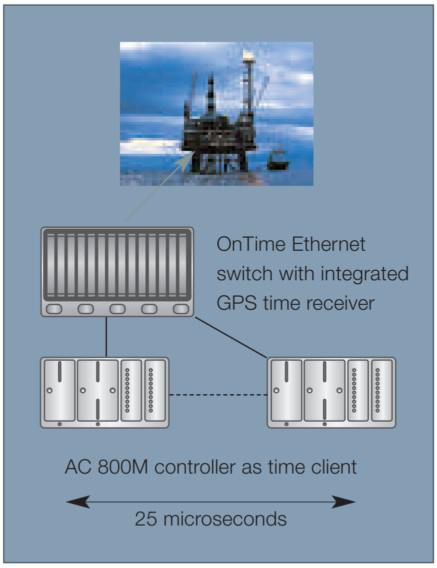

- Full compliance with the SNTP standard
- 0.2 ppm local clock accuracy
- Integrated GPS receiver for system clock generation
- No loss of GPS coverage

An incoming SNTP request packet is time-stamped in hardware as soon as the packet enters the switch, and the corresponding SNTP reply packet is sent when the actual time is equal to the transmit time stamp given in the SNTP payload. The traditional problem caused by non-deterministic access to the Ethernet is not a problem here due to the good interaction between the SNTP timeserver and the switch architecture.

With this time synchronization scheme:

- Timing synchronization accuracy is better than 1 µs when time stamping in the time client is performed in hardware.
- Both server time stamps – T2 (receive) and T3 (transmit) – may be used at the time client for synchronization purposes; plus the drop link propagation delay can be calculated from the round trip time.
- The timing accuracy is independent of the network load.
- No elaborate filtering/erasure techniques are needed in the time client.

The OnTime time synchronization scheme results in an accuracy better than 1 µs. This has been measured using an OnTime switch with integrated SNTP timeserver and an ABB time client based on an ARM7TDMI processor. All time stamping was performed in hardware [3].

This OnTime switch also accommodates less demanding clients. A T3 client, say, directly connected to an OnTime switch with an integrated SNTP timeserver, will fulfill T3 requirements even if the time stamping of outgoing and incoming SNTP packets is performed in an ISR outside the RTOS.

The Grane oilfield installation

In 2002 the class T3 solution discussed in this article was implemented in ABB's Industrial[IT] control system framework as the next-generation standard for high-precision time synchronization. Later in the same year, it was included in a product delivery to the Grane oilfield installation [4]. SNTP client functionality is implemented in the Atlas 044 controller software for the AC 800M controller. OnTime Networks, a collaborating partner, delivered the SNTP timeserver.

The Grane implementation is single-point-of-failure tolerant and will automatically adapt to faulty timeservers and network components. Extensive testing shows a synchronization accuracy of 25 µs between the AC 800M controllers – well inside the Grane requirement. For this project it was of key importance that the ABB solution rely on a standardized protocol and can show absolute time relative to a GPS source.

Looking beyond the success of the Grane implementation, this new Industrial[IT] synchronization concept is also highly relevant for other ABB companies relying on the same kind of functionality.

The future

The solutions presented here for achieving T5 and T3 synchronization standards, with an Ethernet switch from OnTime Networks being used for the former and a standard switch for the latter, have added momentum to the relentless migration towards Ethernet-based standardized time synchronization. Thanks to the allure of significantly reduced cabling, much lower equipment and maintenance costs, and use of low-level time stamp implementation of SNTP, it is safe to say that the move to Ethernet is now inevitable.

Tor Skeie
Svein Johannessen
Trond Løkstad
ABB Corporate Research
Norway
tor.skeie@no.abb.com

Øyvind Holmeide
OnTime Networks AS
Norway

[1] **Ø. Holmeide, T. Skeie:** VoIP drives realtime Ethernet. Industrial Ethernet Book, vol 5, March 2001.
[2] **T. Skeie, S. Johannessen, Ø. Holmeide:** The road to an end-to-end deterministic Ethernet. Proceedings of 4th IEEE International Workshop on Factory Communication Systems (WFCS '02), September, 2002.
[3] **T. Skeie, S. Johannessen, C. Brunner:** Ethernet in substation automation. IEEE Control Systems Magazine, 22(3): 43–51, June 2002.
[4] **D. L. Mills:** Internet time synchronization: the network time protocol. IEEE Trans Communications COM-39, 1482–1493, 1992.
[5] **T. Skeie, S. Johannessen, Ø. Holmeide:** Highly accurate time synchronization over switched Ethernet. Proceedings of 8th IEEE Conference on Emerging Technologies and Factory Automation (ETFA '01), 195–204, 2001.
[6] OnTime Networks, http://www.ontimenet.com/