

The Second Edition of the High Precision Clock Synchronization Protocol

Hans Weibel (hans.weibel@zhaw.ch)
Sven Meier (sven.meier@zhsw.ch)
Zurich University of Applied Sciences
Institute of Embedded Systems (InES)
Technikumstrasse 9
CH-8401 Winterthur, Switzerland

Abstract

A high precision time base is important in many distributed systems. The Precision Time Protocol (PTP) specified in IEEE 1588 is able to synchronize networked clocks with an accuracy down to the nanosecond range. The mechanism combines high accuracy and fast convergence with low demand on clocks and on network and computing resources.

A second generation of PTP has been developed and was approved in March 2008. It was published as standard document IEEE 1588 – 2008, also known as PTP version 2 (PTPv2). This new protocol offers some additional interesting features and improvements which open the door to new applications.

Stellaris® Family of ARM® Cortex™-M3 microcontrollers features an integrated Ethernet interface augmented with hardware assisted IEEE 1588 capability. It is explained how a PTP implementation makes use of the MCU's specific resources.

1 IEEE 1588 Time Synchronization Protocol Overview

IEEE 1588 [1] defines the Precision Time Protocol (PTP) which enables precise synchronization of clocks via packet networks. Compared with alternative synchronization mechanisms such as NTP, GPS, or IRIG, IEEE 1588 provides high precision combined with easy installation. Synchronization and data transfer use the same standard network. No expensive extra cabling or line of sight to satellites is required.

PTP is applied in very different areas.

- Test and measurement: In test and measurement systems, data is acquired by polling the sensors. Sampling timing heavily depends on application program timing and communication latency. A more flexible approach is to equip sensors with a synchronized clock. Each sampled value can then be timestamped for later analy-

sis, or the clock can be used for time-triggered sampling. The LAN eXtensions for Instrumentation (LXI) consortium [2] specifies an instrumentation platform based on industry standard Ethernet technology. LXI makes use of IEEE 1588 to allow triggering directly over the LAN.

- Industrial automation: Ethernet is going to replace field buses more and more. The lack of determinism is compensated with some protocol extensions. IEEE 1588 plays an important role to coordinate communications and actions, and to decouple communication from execution. The presence of precise time information in every device enables distributed synchronous processes to be realized in control applications or in all kind of machinery.
- Power industry: In power plants and substations, voltage and current sensors are used to control and protect the equipment. Event timestamping and data correlation

facilitates applications including fault localization, network disturbance analysis, and detailed recording of events (exact sequence of events facilitates diagnosis). Synchronized sampling, event timestamping, and other advanced functions require precise synchronization. Traditionally a separate synchronization network is used. Being able to transmit synchronization and data over the same network is a big advantage and saves a lot of cabling.

- Telecommunications: In telecommunication networks service quality depends on accurate synchronization. Such networks are traditionally circuit switched and allow the distribution of clock signals over the physical layer. While the networks migrate more and more to packet switching, many traditional circuit switched services continue to exist. An important telecommunication technology is wireless networks. In cellular networks, the handover capability requires precise synchronization of all base transceiver stations.
- Aerospace, navigation, and positioning: In telemetry, radar, and sonar systems, synchronized clocks are generally of importance. Since GPS is expected to be jammed or spoofed in conflict situations, some independence from GPS may be advantageous.
- Audio and video networks: Low latency audio and video transmission over Ethernet enables new applications in residential and studio applications. The IEEE 802.1 working group Audio and Video Bridging (AVB) [3] works on this topic. Synchronization of endpoints and bridges is an important building block of the solution.

2 IEEE 1588 Operational Principle

2.1 Clock Synchronization

The synchronization mechanism is based on a master/slave protocol. PTP instances exchange messages in order to determine the offset between master and slave clocks but also the message transit delay through the network (see figure 1).

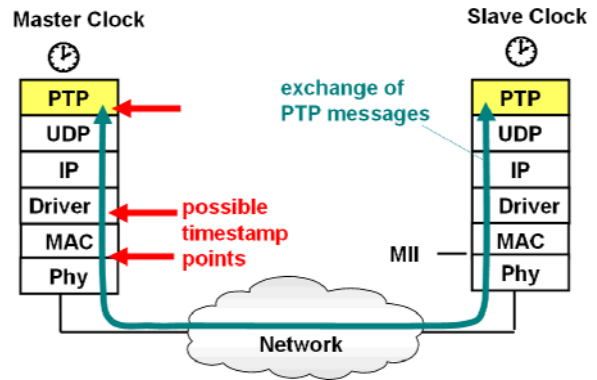


Figure 1: PTP message exchange

Two procedures take place in parallel:

(a) The first task, called synchronization, is responsible to run the slave clock at the same speed as the master. This is achieved by sending a continuous flow of Sync messages from master to slave. Send time t_1 and receive time t_2 of these Sync messages are measured with the local clocks and processed by the slave. The slave's clock has to be adjusted until time intervals are equal on both clocks. Because oscillators are susceptible to environmental changes, Sync messages are sent continuously at a constant rate of typically one or a few messages per second.

There are two options to transport timestamp t_1 to the slave: two-step clocks use the separate Follow_Up message, while one-step clocks deliver t_1 with the Sync message itself. This option requires that the master is capable to insert t_1 into the Sync message on the fly.

(b) The second task determines the slave's offset from the master, i.e. the difference of time of day between master and slave. This is achieved by measuring the two-way delay (round trip time). For the downlink, t_1 and t_2 are available from the last Sync. The uplink is measured with a Delay_Req message providing timestamps t_3 and t_4 (see figure 2).

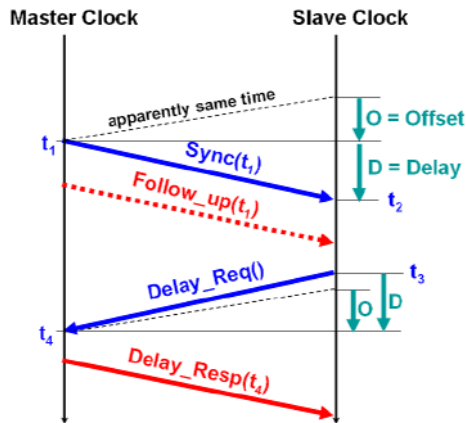


Figure 2: Offset and Delay Measurement

The Delay_Resp message is used to bring t_4 to the slave. Under the assumption of a symmetric transmission path for Sync and Delay_Req, one-way delay and offset from master are computed according to:

$$\text{Delay} = [(t_2 - t_1) + (t_4 - t_3)] / 2$$

$$\text{Offset} = [(t_2 - t_1) - (t_4 - t_3)] / 2$$

Since environmental conditions can change, a continuous correction of the slave clock is required. For this purpose the slave is controlled by a servo loop.

PTP time is represented with 48 bits for seconds and 32 bits for nanoseconds. Where timestamps are transmitted between clocks, an additional correction field extends the resolution to 2^{-16} ns (i.e. 15 femto seconds), paving the way to sub-nanosecond precision. The correction field not only improves the timestamp resolution but also supports the concept for an IEEE-1588-aware bridge type, the Transparent Clock (TC).

2.2 PTP Communication and Network

PTP messages are sent to reserved multicast addresses. Therefore PTP clocks do not need an individual IP configuration.

When the delay of the Sync path varies due to queuing in bridges, the individual measurement results are not very useful.

One approach to overcome this problem is the concept of IEEE-1588-aware bridges. Two network element types exist for this purpose: the Boundary Clocks (BC) and the Transparent Clock (TC).

2.2.1 Boundary Clock

A BC is a bridge equipped with a PTP clock synchronized by the master over one of its ports (see figure 3). Over the other ports, the BC synchronizes slave clocks attached to it.

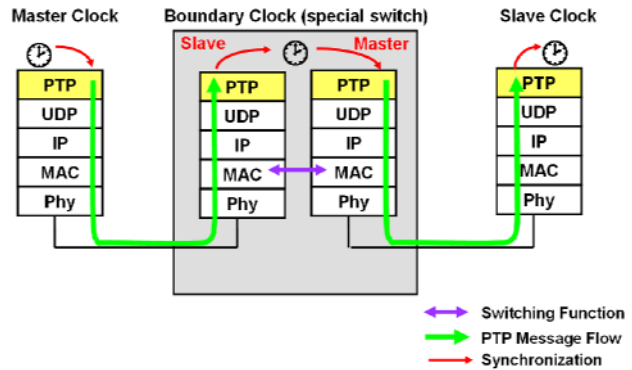


Figure 3: Boundary Clock

Such a configuration represents a synchronization hierarchy (see figure 4) which is established automatically by the so called Best Master Clock (BMC) algorithm. This algorithm takes clock quality and priority settings into account and guarantees that the best available master, the Grandmaster, is the root of the synchronization tree.

The information required to run the BMC algorithm is communicated by Announce messages.

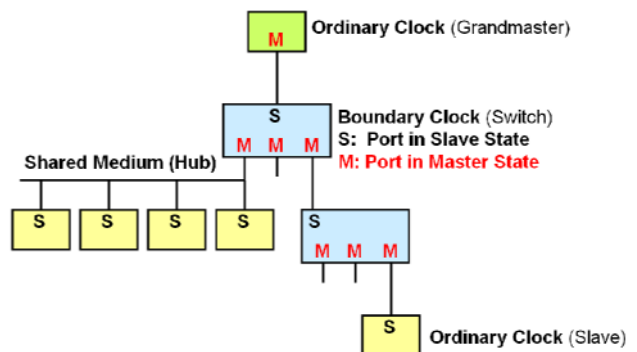


Figure 4: Network topology and clock types

2.2.2 Transparent Clock

The physical layout of a machine determines the topology of an automation network, which is in many cases a daisy chain. When such a topology is built up with BCs, the result is a chain of control loops which is

susceptible to error accumulation. That's why the automation community has proposed the new clock type TC. This is an Ethernet bridge which is capable to measure the residence time of PTP event messages, i.e. the time the message has spent in the bridge during transit. The residence time of the traversed TCs is summed up in the correction field of the Sync message, if the TC is capable to modify the correction field on the fly, or in the respective Follow_Up message. TCs come in two flavors:

(a) In the case of end-to-end (e2e) TCs, the slave measures the delay to the master with an end-to-end delay request / delay response message exchange as described in section 2.1.

(b) The peer-to-peer (p2p) TCs measure the link delay to all neighboring clocks with Pdelay_Req, Pdelay_Resp, and eventually Pdelay_Follow_Up messages. When a Sync traverses a p2p TC, not only the residence time is added to the correction field but also the uplink delay, i.e. the delay of the link over which the Sync has been received (see figure 5).

adjustment while the protocol is implemented in software.

Synchronization accuracy directly depends on timestamp accuracy. The most accurate method is to detect and timestamp PTP messages with hardware assistance as near as possible to the physical layer. For this purpose, transmit and receive data paths are tapped or intercepted in order to capture, decode, and timestamp PTP frames. The required logic can be located in an FPGA, in the PHY, or be part of a microcontroller.

The timestamps are then delivered in some way to the protocol software, together with a message fingerprint (address, sequence number) in order to correlate the timestamp and the respective message.

A slave clock needs to be tunable. It will be accelerated or slowed down according the measurements to reproduce the speed of the master clock as precise as possible.

The protocol software sends and receives PTP messages, fetches the respective timestamps, and carries out the calculations and corrective measures.

More implementation details and hints can be found in [4].

Stellaris® Family of ARM® Cortex™-M3 microcontrollers features an integrated Ethernet interface augmented with hardware assisted IEEE 1588 capability. In order to generate accurate timestamps on the MCU, a mix of hardware and software is used.

First the hardware provides a systick timer with adjustable interval. The interval is adjustable in number of clock cycles (20 ns with 50 MHz clock). The systick timer fires an interrupt every millisecond with the highest priority (nested interrupts) and restarts the timer. Every interrupt another 1'000'000 ns get added to the software maintained clock (sub-millisecond resolution is implemented in hardware, everything else in software).

For timestamping the send and receive interrupt lines of the integrated MAC are connected to two snapshot counters of the CPU.

If a PTP frame (b1 being the size in number of bytes) is completely received, the MAC fires an interrupt. This interrupt causes the counter to take a snapshot of the current counter value c1.




Figure 5: Sync message traversing two peer-to-peer TCs

Since the link delay is measured over all links, even over links blocked by redundancy protocols like Rapid Spanning Tree Protocol (RSTP), a network reconfiguration is seamless with respect to synchronization. No new delay measurement is required when the Sync path changes. A Sync message always reports its own delay, independent of the path it has passed through.

3 Implementing PTP

High precision requires hardware assistance for timestamp generation and clock

In software the current value of PTP time $t1$, the current value of the systick timer $s1$, and the current value of the counter $c2$ get read simultaneously. The difference of the current value of the counter $c2$ and the snapshot taken when the interrupt was fired ($c1$) are subtracted from the current time ($t1 + s1 * 20ns$).

The RX timestamp at the reference point, i.e. at the frame's start frame delimiter (SFD), is calculated as follows:

```
current_time = t1 + s1*20ns
diff_between_int_and_sfd = b1 * 8 * 10ns
diff_between_int_and_now = (c2 - c1) * 20ns
```

```
timestamp = current_time -
             diff_between_int_and_now -
             diff_between_int_and_sfd
```

TX timestamps are calculated accordingly.

To adjust the clock's offset, the systick interval has to be adjusted by adding or subtracting the offset. The offset has to be divided into clock cycles by dividing by the clocks period. This number of clock cycles is evenly corrected over the time of the sync interval to avoid jumps in time. So for a short period of time the interval between two systick interrupts get longer or shorter (depending if negative or positive offset).

Similar to the offset correction the drift gets adjusted by adding or subtracting extra clock cycles to the systick interval.

To use the synchronized time for external event timestamping the same principle as with PTP frame timestamping can be applied (there are 2 more timers left). To trigger events with PTP time the systick interval can be used. This leads to a trigger resolution of 1ms.

Luminary micro provides a driver library to access all the hardware of MCU. It is easy to install and gives full access to all functions provided.

For IEEE 1588 the Institute of Embedded Systems from Zurich University of Applied Sciences provides a full featured PTPv2 stack which has an API to access the clock and to manage the stack. This stack was ported to the LM3S8962 evaluation board and the binary image can be downloaded free of charge [5].




Figure 6: Pulse Per Second comparison between master and slave clock proves an accuracy of ± 100 ns

4 Conclusions

PTPv2 standardized in IEEE 1588-2008 offers a rich set of features, enables a broad range of application scenarios, and fulfills a variety of different requirements.

IEEE-1588-enabled Stellaris® Family of ARM® Cortex™-M3 microcontrollers are equipped with internal resources allowing the implementation of high accuracy PTP clocks without any additional components.

References

- [1] IEEE Std 1588™-2008: "IEEE Standard for a precision Clock Synchronization Protocol for Networked Measurement and Control Systems"
- [2] <http://www.lxistandard.org/>
- [3] http://www.ieee802.org/1/pages/avbri_dges.html
- [4] <http://ines.zhaw.ch/ieee1588/>
- [5] <http://www.ines.zhaw.ch/en/engineering/ines/ieee-1588/software/ptp-for-luminary-cortex.html>