

Introduction

This Application Note is one of a series addressing different aspects of an emerging networking usage model for wireless infrastructure networking equipment: Low Time-Error nodes. The transmission of time / phase information over the wireless infrastructure network with minimal error from the Grand Master to every node in the network is becoming increasingly important for most network operators as they prepare their wireless infrastructure networks to support 5th Generation (5G) wireless protocols and beyond.

Please consult any or all the following documents for a full picture:

- AN-1031 – Time Alignment Background in Wireless Infrastructure
 - Discusses the importance of time-alignment in wireless infrastructure networks and key topics at the network level. Describes the various classes of alignment discussed in ITU-T document.
- AN-1032 – Time-of-Day Transfer within an Ideal Chassis-Based System (this document)
 - Discusses topics related to how to move ToD information from a master timer to various slave timers within a typical wireless infrastructure networking system. This Application Note assumes an ideal system with no wiring or silicon propagation delays on any of the transmission paths. It is necessary to understand the flow of information needed and errors that can be introduced in reading out and loading ToD information into timers. Ensure this note is read and understood before AN-1033 is read.
- AN-1033 – Delay Variation Measurement & Compensation in Non-Ideal Chassis-Based Systems
 - Expands on the discussion in AN-1032 by adding propagation delay effects and how to counter them into the discussion. While using a chassis-based system as an example, single-board systems will experience the same issues and solutions, albeit to a much lesser degree.
- AN-1034 – Minimizing Backplane Signal Usage in Chassis-Based Systems implementing low-cTE Functions
 - This Application Note proposes a method to minimize the number of signal traces needed to transport the necessary information across a system backplane for transfer of accurate ToD from master timer to slave timers. This method allows existing backplanes to be used without extra traces being required. While the method herein could be used in single-board systems, it is not usually as important to limit the number of traces in such a system as it is in a chassis.
- AN-1035 – 8A340x1 ClockMatrix Device Internal Delays and Delay Variations for Compensation Calculations
 - Calculation of exact compensation values is described in AN-1033. This Application Note provides measured values for IDT's 8A340x1 devices for use in those calculations.

Some other related Application Notes that can be of interest are:

- AN-1010 – ClockMatrix Time-to-Digital Converter (TDC)
 - This Application Note describes how to use the TDC circuit in IDT's ClockMatrix family of devices as a precision phase measurement function. Includes details on setups using IDT's Timing Commander software.
- AN-1030 – Input/Input-to-Output/Output Phase Adjustments
 - This Application Note describes how phase relationships can be adjusted input-input, input-output and output-output within IDT's ClockMatrix family of devices. Includes details on setups using IDT's Timing Commander software.
- PWM User Guide
 - This User Guide describes how to use the Pulse-Width Modulation features of IDT's ClockMatrix family of products.
- AN-1036 Using GPIOs for Loading and Latching ToD in 8A3xxxx Devices
 - This Application Note describes the methods for Loading and Latching ToD counters using GPIO signals.

What Happens Within a Node in Terms of ToD Alignment?

In AN-1031, the discussion centered on achieving alignment between adjacent nodes for the master ToD value within each node. However, network nodes typically have many individual ToD counters at different locations within them. For example, PTP algorithms achieve the greatest accuracy when the exchanged time stamp packets can be stamped with the local ToD value (on receipt) or have the local ToD value inserted (on transmit) as close to the physical interconnect as possible. In larger nodes with separate line cards and timing cards, this will be implemented with Time-Stamp Units in each PHY each of which contains its own ToD counter. The network nodes will also have a redundant pair of master ToD counters on two separate timing cards. In small systems the peripheral ToD counters can be located more centrally, such as in a switch fabric because the delay variation is much smaller in those systems and they may not support redundant master ToD counters.

Figure 1 shows an implementation of ToD synchronization within a router node which in this case has a redundant pair of timing cards and 16 line cards, each with its own time-stamping PHY device (with internal ToD counter). Not all extra copies of connections are shown for simplicity.

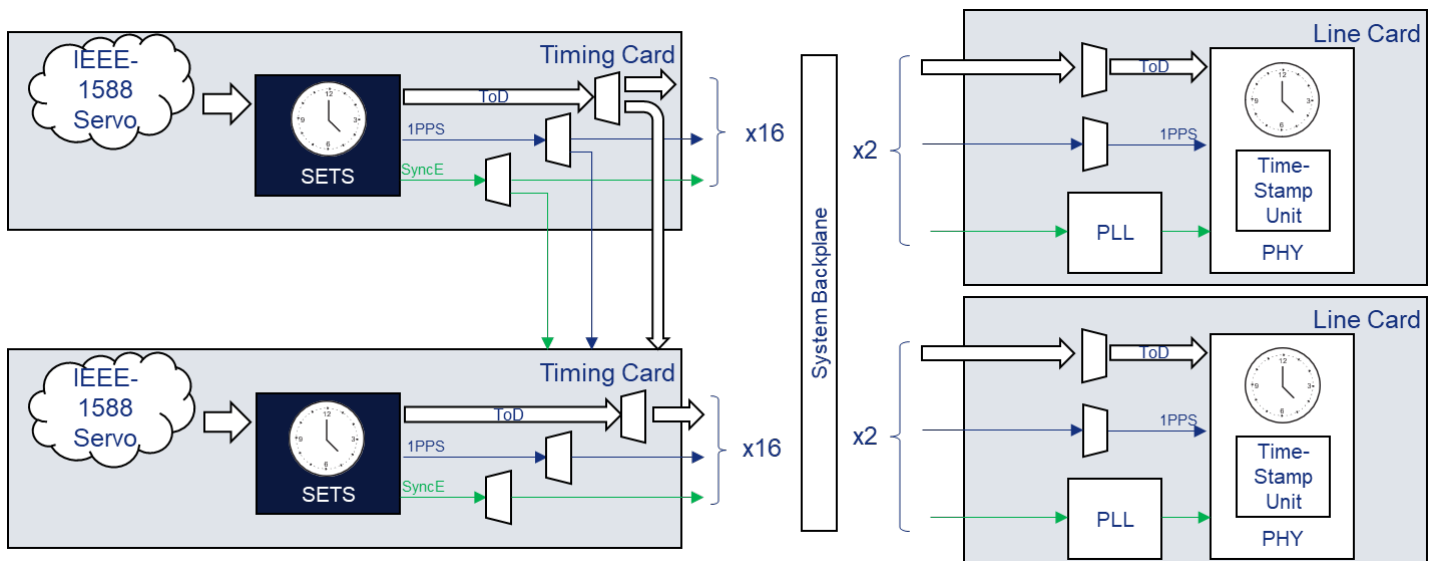
How Do Base-Stations Establish and Maintain Alignment?

Actual base-station specifications established by 3GPP are about the phase alignment of radio signals. This can be established within each base-station as a Time-of-Day or as a phase alignment.

One method that has been used is to have a Global Navigation Satellite System (GNSS) receiver at each base-station. This works well if a receiver can view enough of the sky to be in contact with 4 of the (usually) 8 GNSS satellites that are above the visible horizon at any time. Base-stations so equipped can achieve an alignment of $\sim \pm 50\text{ns}$ under optimal conditions. Tighter alignments are possible, but only with more expensive techniques such as differential GNSS. Other issues with this method include interference in the satellite signals as well as Line-of-Sight issues that prevent enough view of the sky.

As shown in AN-1031, it is also possible to transmit time/phase information from a central time source in a network to all nodes in the network over the network inter-connect. This removes the Line-of-Sight and signal interference issues of GNSS but is subject to a great deal of both fixed and varying delays.

Figure 1. Implementation of ToD Synchronization with a Router Node



To transfer ToD from the active timing card and maintain its alignment once transferred, there are 3 elements required:

- A ToD value. This is loaded at the line card or redundant timing card. This is a piece of data, often in a standard representation. This paper will use the 88-bit format shown in Figure 2 in examples, although the format can be anything that devices at either end of the connection can parse. Transmission of this data is not time-critical if it arrives before the relevant synchronization pulse.
- A synchronization pulse. This tells the destination device precisely when to load the ToD data into its internal counter and start counting with it. In the examples herein, this will be shown as a 1PPS signal. Other frequencies can be used, but they must be slow enough to allow the ToD information transfer including any decoding or setup time. So 1PPS is a common choice. This is usually implemented as a repetitive, periodic train of pulses, rather than just a one-time pulse.
- A master clock signal. This is a common clock used by all ToD counters so that once loaded they will count synchronously and so not drift apart between updates. The phase alignment of this signal from one card to another is not critical, only its frequency. Most T-BC systems will be on a Synchronous Ethernet (SyncE) network for frequency stability and the recovered SyncE clock is often used for this purpose. That is what is illustrated in Figure 1.
 - There is one warning in this statement: if the master clock is a factor in how quickly the ToD value is loaded when a synchronization pulse is received or how quickly it can be read for use in the time-stamp unit (i.e. synchronous load / read), then it can introduce a quantization error that can affect the overall cTE value. If this is the case, a higher frequency will result in less error contribution.

Achieving Accurate ToD Alignment Within a Node

As discussed in the previous section, ToD information needs to be transferred from the master ToD counter to the peripheral ToD counters on the backup timing card and on the line cards. This requires transmission of both ToD data and a synchronization pulse from the active timing card to each destination. Since the delays in transmission from master ToD to each peripheral ToD within the system are different, it is necessary to compensate for them, either by adjusting the timing of the synchronization pulse to each line card (and potentially to each destination within each line card) or by pre-compensating the ToD value that is transferred to each destination. Both methods will be discussed in this document.

Relationship Between Synchronization Pulses, Master Clock Signals and ToD Values

In PTP systems, ToD Values are 88-bit data structures consisting of 3 fields as shown in Figure 2.

Figure 2. 88-Bit Data Structure Consisting of 3 Fields



This value represents the elapsed time since a known time reference referred to as an epoch. For example, the 1588 Epoch is from January 1st, 1970. A ToD value reference to the 1588 Epoch would represent the number of whole seconds, nanoseconds and fractions of a nanosecond since that point in time. Note that the number of nanoseconds rolls-over after 999,999,999nsec since a nanosecond is one-billionth of a second. This does not use up the full range of values in the 32-bit nanoseconds field.

ToD counters (sometimes called accumulators) are loaded with a starting value with a synchronization pulse and then count upwards using a master clock. Each period of the master clock, a time value representing the period of the master clock is added to the counter. For example, a 100MHz clock has a 10nsec period, so at each rising edge of the 100MHz clock, 10nsec are added to the ToD value. In between rising edges, the ToD value is not updated. This behavior implies that to minimize ToD errors, certain relationships between the synchronization pulses, ToD values and master clock signal need to be maintained.

Figure 3 shows a logical schematic of a ToD counter within an integrated circuit. The synchronization pulse (SYNC IN) that is used to load the ToD counter arrives at Point A, which is the pin of the integrated circuit. It will experience some internal path delays before arriving at the counter itself at Point C. Similarly, the master clock will experience delays from the pin at Point B to the counter at Point C. These two path delays will be different by small amounts. The actual ToD data to be loaded into the ToD counter also has its own path to arrive at the counter internal to the integrated circuit. However, if the ToD data is there more than the required setup time before the counter loads, those delays are not relevant.

Assuming the ToD counter is a synchronous loading counter (i.e. loads on the rising edge of its clock input, not on the rising edge of the load signal (LD)), and if the Point A to Point C delay is much less than one period of the master clock, and the synchronization pulse is synchronous with the master clock, those delays also don't matter. However, if the delay from Point A to Point C is longer than one clock period or the synchronization pulse is not synchronous with the master clock, it is impossible to be sure which rising edge of the master clock will latch the ToD data into the counter. This can introduce an uncertainty of one or more clock periods of the master clock into the ToD value that is loaded.

Figure 3. ToD Counter Logic Diagram

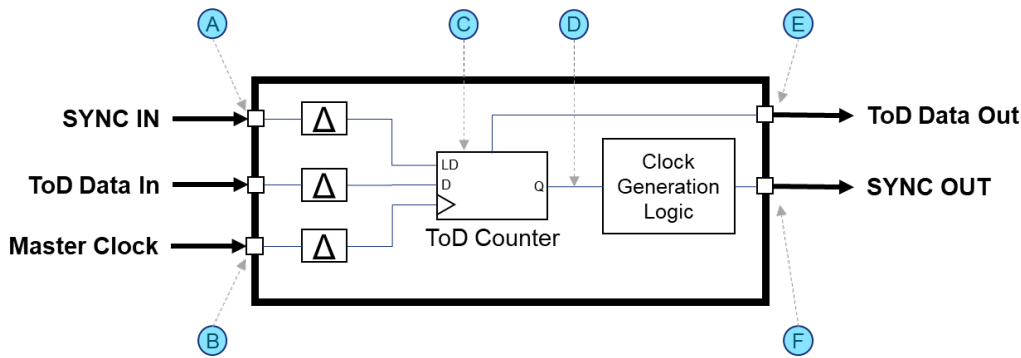



Figure 4. ToD Discrepancy When Loading a ToD Counter



The recommended (synchronous-load counter) case is illustrated in Figure 4. This diagram assumes the ToD counter loads at the rising edge of the first master clock edge after the load pulse, where the load pulse is aligned with the desired (actual) time.

The accuracy of the ToD loading as shown in Figure 4 is driven by two factors:

- The relationship between the incoming synchronization pulse and master clock. This offset is fixed. It is also visible since it is between two signals that are external to any integrated circuit, so it can be measured.
- The delay internal to the integrated circuit from Point B to Point C. This delay is internal and so can't be measured directly but can be provided in the integrated circuit's datasheet. The delay will be largely fixed with a small variation due to temperature and/or voltage changes.

By using the measured and known values of the fixed portion of the two delay factors, a compensation factor can be applied to the ToD value loaded into the ToD counter. In the example above, if a ToD value of 1000.000000003.0 were loaded, then the internal ToD counter would be synchronized with the actual ToD.

An internal ToD counter can also be used to generate time-aligned synchronization pulses. An output synchronization pulse (SYNC OUT), can be generated with rising edges (and/or falling edges) aligned to specific time intervals. For example, to generate a 1PPS output, every time the ToD counter rolls-over the nanosecond sub-field, it will increment the seconds sub-field and generate a rising edge on the clock. The falling edge of that clock could be synchronized to an intermediate count or just stay high for an approximate pulse width (e.g. 10msec) before returning low. Note that the logic to increment the ToD values is error-free because the synchronous counter's combinatorial logic performs its work in the interval between master clock edges and is correctly loaded at the rising edge. The same is not true for the synchronization pulse as shown in Figure 3 because there is at least some logic delay in generating the rising edge (Point C to Point D) and some further delay (Point D to Point E) in propagating this signal to an external pin outside the integrated circuit. Not only will this delay likely be several nanoseconds, it will also vary with silicon performance parameters such as Process, Voltage or Temperature (PVT). The synchronization pulse will have an error that can be significant and is not known under all conditions and so can't be compensated for elsewhere in the system.

In Figure 5, the rising edge of the 1PPS synchronization pulse, which is supposed to signal a 1-second roll-over point of the ToD value is inaccurate by anywhere from 1 nanosecond up to 6 nanoseconds. This could be treated as a fixed error of 3.5nsec (which can be compensated for) with a variable error of ± 2.5 nsec (which can't be compensated for).

Figure 5. Generation of a Sync Out Pulse from a ToD Counter

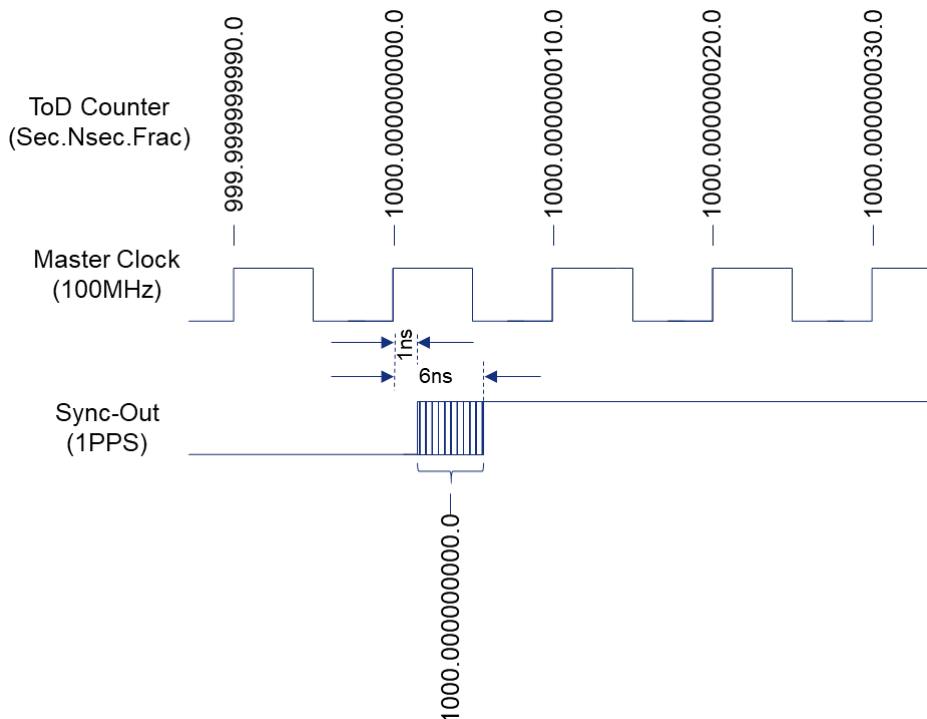


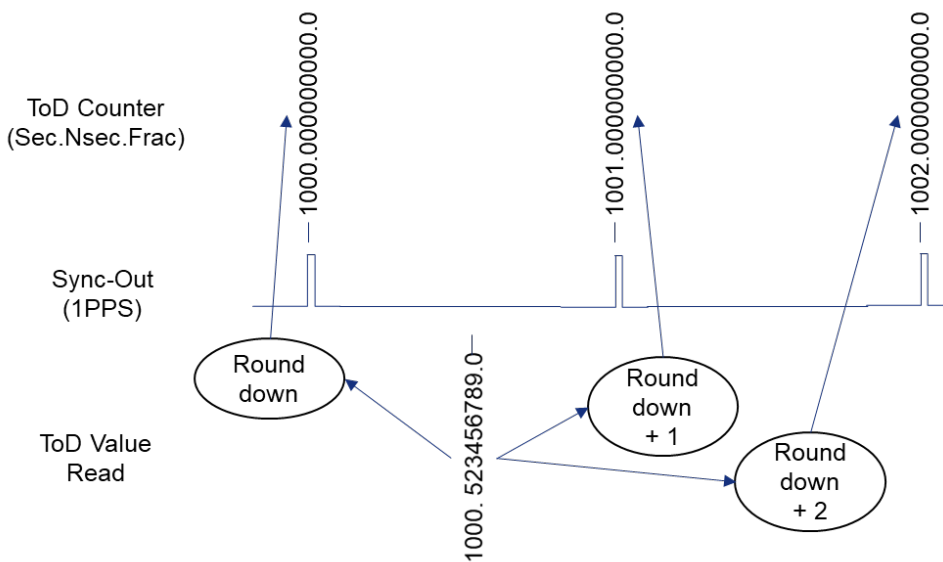
Figure 5 uses a 1 second roll-over point to generate a 1PPS synchronization pulse, but different intervals can be used, and different pulse rates can be generated. For example, a 1kHz rate could be generated by logic that triggers when the nanosecond portion passes through a multiple of 1,000,000.

In addition to loading a value into a ToD counter and to generating a synchronization pulse, there is also the need to read the ToD value accurately to pass to other ToD counters elsewhere in the system. When a ToD value is read, it represents the ToD contents at the point in time where the read signal arrived at that counter. Since the timing of a read signal internal to an FPGA or integrated circuit is often unknown and usually highly variable it is best to correlate ToD values with the synchronization pulses.

Since we know the synchronization pulse was generated at a known roll-over point of the ToD counter we are reading, we can round down the value we read from the ToD counter modulo the synchronization pulse interval. This will give us ToD value associated with that synchronization pulse. For example, if we are generating synchronization pulses every 1 second (1PPS), then we know the nanosecond and sub-nanosecond fields of the ToD were zero at that time. By simply taking the seconds field value, the ToD value at the previous synchronization pulse is known. Furthermore, what the ToD value would be at any future synchronization pulse can be extrapolated by adding an integral number of seconds to that rounded ToD corresponding to the number of synchronization pulses that we want to pre-compensate for.

Figure 6 illustrates this using a 1 second interval.

Figure 6. Reading, Rounding and Extrapolating ToD Values

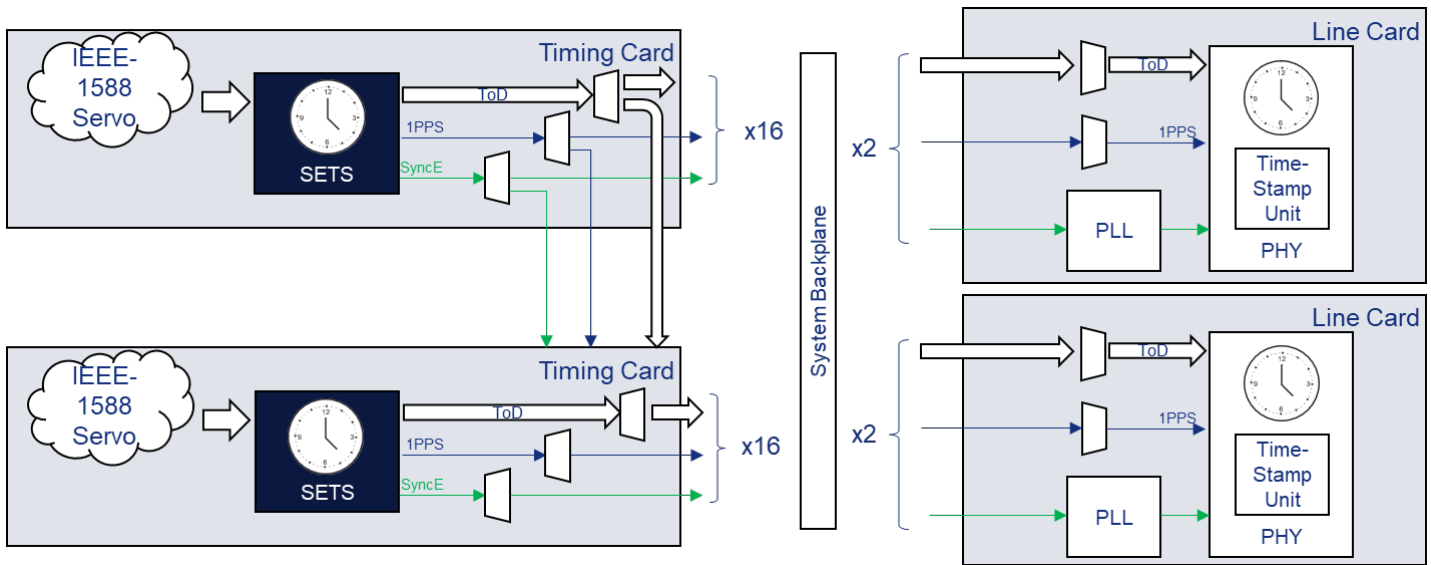


The above method of extrapolation does take some time to calculate and further time to distribute the desired ToD values around a system. For that reason, the synchronization pulses must have a large enough interval to allow for reading from the master ToD, calculating the ToD value to use for the peripheral ToD, transporting the ToD to the peripheral ToD location and pre-loading it into the peripheral ToD counter, ready to be synchronously loaded at the next synchronization pulse.

In practice it has been found that a synchronization pulse rate of less than 1kHz is needed and quite often 1PPS is used for convenience.

Putting This Together in an Ideal System

Returning to the system overview reprinted in Figure 7, the active timing card (shown at upper left in the diagram) contains the master ToD counter in the device labelled SETS. The master ToD value has been set by the IEEE 1588 Servo function using T-BC protocols with the upstream system. In Figure 7, that master ToD is generating synchronization pulses at a 1PPS rate. The 1PPS pulses are aligned to the 1 second roll-overs of the master ToD counter with a small uncertainty as described earlier.

Figure 7. Implementation of ToD Synchronization with a Router Node


The ToD value is read from the master ToD, then rounded down to the nearest second (or whatever the period of the synchronization pulse is). Since this is an ideal system, the ToD transport path is fast enough to get data to all peripheral ToD counters before the next synchronization pulse rising edge. One synchronization pulse period (1 second in these examples) is added to the rounded-down ToD value at the master which is then transported over the ToD data path and pre-loaded in all the peripheral ToD counters. At the next rising edge that pre-loaded value will take effect.

In an ideal system, the synchronization pulses from the master ToD counter would arrive at the peripheral ToD counters in zero time and would unambiguously load the correct ToD value at the same point in time with the same ToD value at all destinations. Then the only error in the peripheral ToD values versus the master value would be the uncertainty in generating the Sync-Out pulses at the master ToD plus the small path delay variation in the master clock path within the peripheral ToDs.

Once the peripheral ToDs are properly loaded, then as long as they are clocked from the same frequency clock as the master ToD they will not diverge from each other. The system designer should ensure all are clocked by the same master clock frequency.

If there were no disruptions in the master ToD value, then it would only be necessary to perform the above process once. However due to network disruptions it is likely that the PTP protocol will occasionally adjust the master ToD value, which then needs to be re-distributed. It is recommended, even in an ideal system, that the process be repeated periodically. How frequently is dependent on network parameters beyond the scope of this discussion.

Summary

In this application note, it is shown that even in an ideal system where there are no delays in transporting signals and data from the master ToD counter to the peripheral ToD counters, there are still several important design considerations to address. Also, there are several contributors to system time error that need to be understood and accounted for. The contributions to system time error from within the integrated circuits containing the master and peripheral ToD counters include both fixed portions, which if known can be compensated for, and variable portions which add directly to the system time error budget. It is essential to use integrated circuits where these delays have been characterized and published by the silicon vendors to achieve the lowest system time error.

Glossary of Terms

Table 1. Glossary of Terms

1PPS	One Pulse-pre-Second – common synchronization clock used within networking systems. Aligns with the 1 second roll-over of the ToD timer that generates it. This is a 1Hz periodic signal but can be referenced to the rising or falling edge of each pulse. A 1PPS pulse is not required to have a 50% / 50% duty cycle.
3GPP	3 rd Generation Protocol Partnership – international standards body that defines specifications for wireless communications.
5G	5 th Generation Wireless Networking family of protocols. This is an imprecisely defined term that is loosely understood to refer to the IMT-2020 series of standards being defined by 3GPP.
CP	Charge Pump – sub-circuit within an analog PLL that converts the time-related pulse width from a Phase Detector into a control voltage that can be applied to the VCO.
cTE	Constant Time Error – error in time introduced by a single network node as it receives and transmits time / phase information. This is specifically the error introduced by fixed functions that cannot be accounted for or compensated for.
DCO	Digitally-Controlled Oscillator - sub-circuit within a PLL that contains an oscillator generating the master frequency reference within the PLL. The frequency of oscillation can be adjusted by applying different digital control values. These values are often called Frequency Control Words.
dTE	Dynamic Time Error – error in time introduced by a single network node as it receives and transmits time / phase information. This is specifically the error introduced by periodically varying factors that cannot be accounted for or compensated for.
FIFO	First-In / First-Out – a circuit that queues up pieces of information and maintains them in the order they were received.
FOM	Fiber-Optic Module
FPGA	Field-Programmable Gate Array – often used by board designers to implement complex circuits. Due to their flexible nature, FPGAs are ideal for implementation of circuits that can need changing to fully achieve the desired functionality. Unfortunately, that flexibility also leads in many cases to large uncertainty in signal propagation delays.
FR4	Fire Retardant 4 – material used in Printed Circuit Board manufacturing. FR4 is one of the cheaper options and is widely used in PCBs that don't require tightly controlled impedances, low parameter variation or support multi-GHz frequencies.
GM	Telecom Grand Master – common time reference source for a wireless infrastructure (or other) network. Capable of transmitting its reference using the IEEE 1588 protocols to other network nodes. Usually only one GM is active in any network, although one or more backups can be available to take over in case of disqualification of the active GM.
GNSS	Global Navigation Satellite System – generic term used for any of several satellite constellations used to transmit time information used to establish position on the globe or used as a common time reference. The protocols used require at least 4 satellites to be visible to any receiver to achieve the time/phase alignments discussed in these Application Notes
IDT	Integrated Device Technology – a semiconductor designer and supplier with a leading market position for timing and synchronization integrated circuits, including the ClockMatrix family.
IoT	Internet-of-Things – term used to include any device of any kind connected to the Internet. Recently there has been an exponential increase in the number of such devices as simpler devices are now able to cheaply connect to the Internet over wireless networks.

ITU	International Telecommunications Union – a standards body that defines internationally recognized specifications for the interaction of telecommunications and networking equipment
LF	Loop Filter – sub-circuit within a PLL that takes the digital words or analog control voltages from the PD and CP sub-circuits over a period of time and filters them to ensure the VCO responds smoothly to the requested changes. This is usually a low-pass filter, although more complex digital filters can be used.
PD	Phase Detector – component within a PLL device that detects the difference in time between the rising (or falling in some cases) edges of the two input signals. The output can be a voltage or pulse that is proportional to the time difference in analog PDs or a digital word in the case of digital PDs (sometimes also called TDCs).
PHY	Physical layer protocol translation device. In the context of these Application Notes, these are assumed to include a Time-of-Day counter used to time-stamp IEEE 1588 packets for minimum inaccuracy in the transfer of time/phase information.
PLL	Phase-Locked Loop
PTP	Precision Time Protocol – any protocol used across a communications network for transferring time information. In the context of these Application Notes, it refers specifically to IEEE 1588
PWM	Pulse-Width Modulation – a way of encoding extra information onto a periodic signal such as a clock
RTT	Round-Trip Time – elapsed time from when a signal is transmitted until it is looped back and received at the source
SerDes	Serializer / Deserializer – circuit within a PHY device that generates the signals on the line side of the PHY with very high speed. Reference clocks for the SerDes usually need to meet low phase noise targets to maintain low bit-error rates on the line (clean eye pattern).
SETS	Synchronous Equipment Timing Source – function within a telecommunications system that establishes and communicates the master time sources for the node
SyncE	Synchronous Ethernet protocol – defined by ITU-T G.8261. Carries data in similar packet formats as regular Ethernet, but also includes a frequency reference that can be used by all nodes in the network. Requires an unbroken path to the active GM of the network across synchronous networking protocols.
T-BC	Telecom Boundary Clock – one of several profiles within the IEEE 1588 family of standards. Defined in ITU-T 8273.2.
TDC	Time-to-Digital Converter – digital circuit that measures and reports time differences (phase errors) between the rising edge of two signals supplied to its input terminals. The output is a time measurement that can be used to drive a phase-locked loop or read by external software for use in compensation calculations. Some TDCs require periodic signals to achieve a desired accuracy whereas others can make one-shot measurements.
ToD	Time-of-Day – in the context of these Application Notes, this does not refer to a universal time standard, but rather to the time representation used in a specific wireless infrastructure network.
UTC	Universal Time Content – internationally recognized accurate time standard and format for representing it.
VCO	Voltage-Controlled Oscillator – sub-circuit within a PLL that contains an oscillator generating the master frequency reference within the PLL. The frequency of oscillation can be adjusted by applying different control voltage values.
ZDB	Zero-Delay Buffer – an application of a PLL to create multiple copies of an incoming clock signal, potentially at different output frequencies, but with minimal delay from input-output. Despite the name, delays from input-output are not zero, but are much smaller than a normal PLL implementation. A ZDB usually uses an external PCB trace to implement its feedback path. This PCB trace acts as a delay function of one period of the clock being fed back.

Revision History

Revision Date	Description of Change
February 8, 2019	Initial release of the application note.

**Corporate Headquarters**

6024 Silver Creek Valley Road
San Jose, CA 95138
www.IDT.com

Sales

1-800-345-7015 or 408-284-8200
Fax: 408-284-2775
www.IDT.com/go/sales

Tech Support

www.IDT.com/go/support

DISCLAIMER Integrated Device Technology, Inc. (IDT) and its affiliated companies (herein referred to as "IDT") reserve the right to modify the products and/or specifications described herein at any time, without notice, at IDT's sole discretion. Performance specifications and operating parameters of the described products are determined in an independent state and are not guaranteed to perform the same way when installed in customer products. The information contained herein is provided without representation or warranty of any kind, whether express or implied, including, but not limited to, the suitability of IDT's products for any particular purpose, an implied warranty of merchantability, or non-infringement of the intellectual property rights of others. This document is presented only as a guide and does not convey any license under intellectual property rights of IDT or any third parties.

IDT's products are not intended for use in applications involving extreme environmental conditions or in life support systems or similar devices where the failure or malfunction of an IDT product can be reasonably expected to significantly affect the health or safety of users. Anyone using an IDT product in such a manner does so at their own risk, absent an express, written agreement by IDT.

Integrated Device Technology, IDT and the IDT logo are trademarks or registered trademarks of IDT and its subsidiaries in the United States and other countries. Other trademarks used herein are the property of IDT or their respective third party owners. For datasheet type definitions and a glossary of common terms, visit www.idt.com/go/glossary. All contents of this document are copyright of Integrated Device Technology, Inc. All rights reserved.