

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
 - 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 (this document)
 - 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.

How Much Delay Variation Can Be Expected in a Non-Ideal System

In AN-1032, issues with transferring ToD information across an ideal system were discussed. The main issue to deal with in an actual system is that not only are the synchronization pulses not transported in zero time to all destinations, the transport delay is different for each destination and can also vary with changes in system voltage or temperature. Furthermore, the delays can be different from one system to the next due to manufacturing tolerance issues.

When the system needs to meet the time error requirements of Class A or Class B, it can be enough to measure the delays statically in a lab or production environment, add appropriate guard-bands to deal with process, voltage or temperature (PVT) variations and use those values to calculate compensation values. This can allow adjustments to ToD values in various peripheral devices to meet those more relaxed specifications.

However, as requirements tighten in the Class C or Class D systems, these methods will not be enough to meet the time error limits for the systems. System temperature changes can cause signal propagation delays in PCB traces to vary by 1% for an 80°C temperature change. Typical router systems assume a 40 inch trace length as a long trace across a backplane. With electrical signals in FR4 material near 170ps/inch, $40\text{inch} * 170\text{ps/inch} * 1\% = 68\text{ps}$.

The rise time of CMOS signals in a heavily loaded backplane environment is also affected by several factors. CMOS signals are often used across a backplane to save pins / traces when transporting low-frequency signals such as synchronization pulses. Based on characterization data on IDT's ClockMatrix devices, CMOS output rise times can vary from 110ps to 510ps due to PVT variations. Assuming the switching point is at the half-way point of the rise, that results in a variation of $510/2 - 110/2 = 250\text{ps}$.

In addition, the propagation delay through a CMOS buffer such as IDT's 5PB1108 can vary as much as 1nsec over PVT.

Taken together, sending a signal from an IDT ClockMatrix device through an 5PB1108 fanout buffer across 40 inches of FR4 material to a receiver on another card could add $250\text{ps} + 1000\text{ps} + 68\text{ps} = 1318\text{ps}$ of variation in delay. That is the amount of guard-band a statically configured system will have to add. This will directly subtract from the system time error budget.

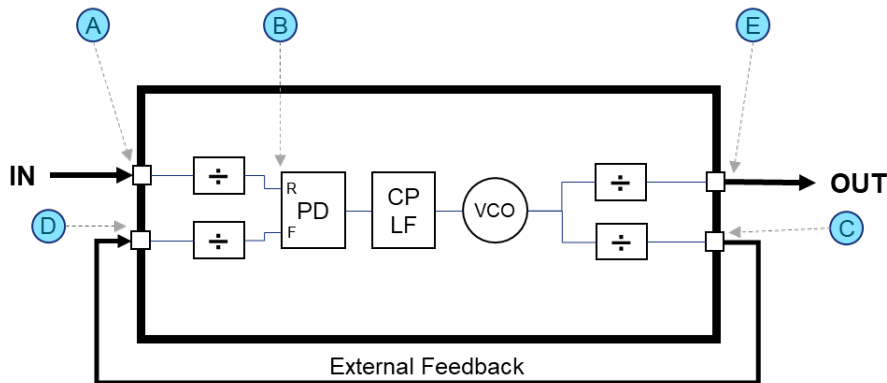
For Class C or better systems, it will be necessary to move to a dynamic, real-time measurement method. By simply measuring the delay in-system once per system reconfiguration, manufacturing process variations are removed from the total variation. Note that it is not enough to measure a system once at power-up. Since modules can be inserted and withdrawn, it is necessary to recalibrate every time the system configuration changes.

By repeating these measurements periodically, the effects of factors like voltage and temperature can also be factored out or at least reduced to within tolerable limits.

There are two techniques applied in dynamically-adjusted systems. The Zero-Delay Buffer (ZDB) technique is used for small portions of the total signal propagation path. The round-trip time (RTT) method is used for the remaining portions of the signal propagation path, or at least the parts that introduce enough variation in the overall total time error.

How a Zero-Delay Buffer (ZDB) Works

For many years, silicon devices called zero-delay buffers have been used to propagate and fan-out clock signals with minimal signal delay. Despite the name, none of these devices introduced zero delay, although they did significantly reduce it compared to standard buffers. The name is also deceptive in another way. These are not buffers, but actually Phase-Locked Loop (PLL) devices.

Figure 1. Phase-Locked (PLL) Used as a Zero-Delay Buffer


Referring to Figure 1, the principle of operation of any PLL circuit is to perfectly align the reference (R) and feedback (F) inputs to its Phase Detector (PD). The PD indicates a positive or negative difference between the rising edges of its inputs. That error value is passed through the Charge Pump (CP) and Loop Filter (LF) stages and will speed-up or slow-down the oscillation frequency of the Voltage-Control Oscillator (VCO). Since the output of the VCO is fed back to the PD, eventually the two edges will align at the PD (Point B in the above diagram), resulting in no error signal and the VCO attaining a steady-state frequency. This is called a Locked condition. Any changes in the edge positions on the input being tracked (IN) will result in the output edge position moving with it.

In a standard PLL, the VCO is fed back to the PD via an internal signal. Since the delays in the IN to OUT path are not equivalent to the delays in the feedback path, the difference in time between rising edges on the IN versus the OUT signals will be large and can also vary significantly from one PLL to another and with changes in voltage or temperature.

In a zero-delay buffer though, rather than an unmatched internal loopback, an external loopback is used. Furthermore, the external loopback signal is delayed by a calculated amount to achieve a small apparent IN – OUT delay and to minimize variation in that delay. In an ideal ZDB, that delay would be one period of the OUT clock.

To minimize variations in delay, the frequencies of IN, OUT and External Feedback must be the same. That way all the dividers shown are configured the same way. Similarly, the signal types used on IN, OUT and External Feedback must also be the same to minimize variations in the input and output buffer behaviors. The two output points (C and E) must be supplied by the same voltage rail and be physically near one another to minimize any voltage or temperature-dependent behavior differences between them. Similarly, the inputs (A and D) must also share voltage and location.

Assuming all those factors are taken care of, there will be minimal difference between the delays from A to B versus D to B and in B to E versus B to C. By delaying the External Feedback signal by one or more periods (T) of the OUT clock, IN and OUT will appear perfectly aligned.

Non-Ideal ZDB Conditions

In the above explanation, it was assumed that the path delays, especially those from A to B versus D to B were the same. By those being the same, then the error determined by the PD at its inputs is the same as the phase difference at the input pins of the device. However, that isn't always the case. Differences in input buffer protocols, trace routing, internal signal loading and small temperature or voltage differences inside the device can all cause the signals to take different amounts of time to move from the pins to the PD.

If that is the case and the differences are of non-trivial magnitude, then these need to be considered or they will affect the output phase alignment. The suppliers of the ZDB device need to provide this information for consideration. The information will list all possible paths through the device from input pins to one or more PDs and to the R or F input of each PD. The important information is the difference in delay of the two paths. If one path is 5.1nsec and the other is 5.2nsec, then difference is 0.1nsec, which is the value that needs to be compensated for.

The measurement will consist of two parts: a fixed portion (or mean value) and a variable portion (or deviation range). The fixed portion of these delays can be compensated for within some ZDB devices, such as IDT's ClockMatrix family. Please see AN-1030 – Input/Input-to-Output/Output Phase Adjustments for details on how to perform this compensation. One method is to subtract the known offset from the error signal the PD produces before that error is passed onto the CP block for processing.

The variable portion will form part of the total error budget, since there is no way to determine what compensation value to apply at a time.

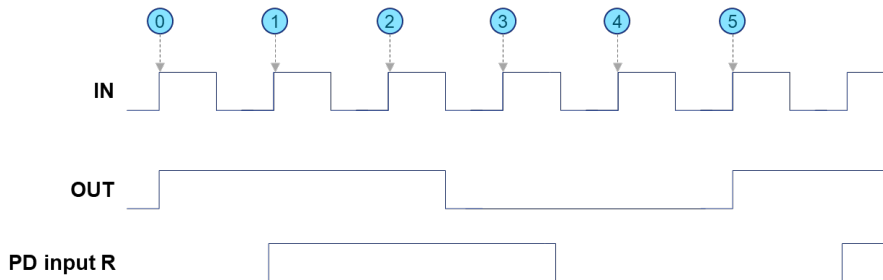
Another issue that can arise is that not all signals involved are the same frequency. If for example, IN and OUT are different frequencies, this can affect the accuracy of the measurement due to different delays on different paths. Note that IN and OUT must be related in such a way as to periodically have their edges coincide or there is no sense in trying to align them. If for example, IN = 125MHz and OUT = 25MHz, then every 5th edge of IN will align with the OUT edges as shown in Figure 2.

In terms of setup in this case, it is recommended that the External Feedback signal use the same frequency as OUT. That way the path delays from B to C and B to E will remain as close as possible. The output protocols (e.g. LVDS, LVPECL, etc.) of OUT and External Loopback should be the same since the output buffer delay variation for different output protocols will be much larger than many other error sources.

Since External Loopback is a different frequency than IN, then the input dividers in their respective paths will be programmed differently to ensure the same frequency signal arrives at the PD. This will result in different delays on those same paths. Note that for the same reason as on the output side, it is recommended that both IN and External Loopback use the same signal protocols. Using the example discussed above, External Loopback would not be divided since it is already 25MHz, but the IN signal would be divided by 5 so that the PD is comparing two 25MHz signals. The divider will add some delay to the A to B path. Information on that delay for specific divider values will be provided by the device vendor and can be dealt with as indicated above for different path delays.

However, there is one more issue to consider. This is the issue of aliasing. As shown in Figure 2, there are 5 rising edges of IN for each rising edge of OUT.

Figure 2. Illustration of Phase Detector Aliasing Issue



The rising edge of OUT is intended to align with edge 0, edge 5, edge 10, etc. of IN. Other alignments will result in different apparent delays in the OUT signal relative to other signals in the system, adding to error budgets. However as shown above, the divider on path A to B is not synchronized with the External Loopback path, so the divider can start its work at IN edge 0 as desired or it can start at edge 1 to 4, resulting in one or more multiples of 8nsec (period of 125MHz) being inserted in the delay of OUT. Once the divider starts-up, this offset will remain constant until some configuration change causes a new relationship to happen (such as loss-of-signal of IN followed by reacquisition of signal).

Compensating for this requires some external information to determine how the divider start point aligns with the desired start point. If OUT is compared to some other signal elsewhere in the system and some fixed error in multiple of 8nsec is observed, then an appropriate compensation value can be added as indicated previously. Since this issue is difficult to detect and compensate for, it is recommended that IN, OUT and External Feedback all use the same frequency.

Another non-ideal behavior is the output-output skew of the External Loopback versus the OUT signal. Most devices provide an output-output skew specification, which could be used directly as the error contribution. However, that is often too pessimistic and with additional information, there is no need to add this full value to the system error budget. Typical device datasheets indicate a maximum skew value X, without indicating direction. The user needs to assume Signal A can lead Signal B by X and under some other condition, B can lead A by X. This leads to a range of possibilities from +X to -X in terms of signal uncertainty. The actual measurements are usually not that bad. In many cases, two signals go through the same series of internal buffers and PCB traces. This results in one signal always leading the other. Variations in temperature, voltage or process will change the magnitude of the lead, but not the fact that one always leads. This isn't always the case since different buffer types can be used on different output paths. This can result in cases where one leads under some conditions and the second leads in other cases. Even in this situation, the magnitude of the edge movement is usually much less than +X.

Also, output-output skew specifications in device datasheets often don't show a minimum value. Not only does one signal usually lead the other, there is usually a non-zero minimum lead time. An uncertainty of $+X$ to $-X$ is likely reduced to $+X$ to $+X/4$ or something similar. Knowing this can reduce the uncertainty by a significant amount.

Detailed information on output-output skews from the device vendor can make this clear and avoid adding excessive uncertainty.

Output-output skews are so defined for IDT's ClockMatrix family (in AN-1035) and many buffers. Also, uncertainty can be compensated for by adjustments within devices like the ClockMatrix family.

The External Loopback signal is usually designed using a calculated PCB trace length to achieve the desired one clock period delay. That calculation is performed using an assumed delay per unit length of PCB trace in the chosen PCB materials and stack-up. For example, 170ps/inch is a common number for FR4 PCB material in a stripline configuration. However, that delay is a nominal value that is affected by temperature and manufacturing tolerances. Compensation values may need to be determined and included in the ZDB path to deal with any non-ideal behavior in this path. Some techniques on how to do this will be presented in the RTT implementation section.

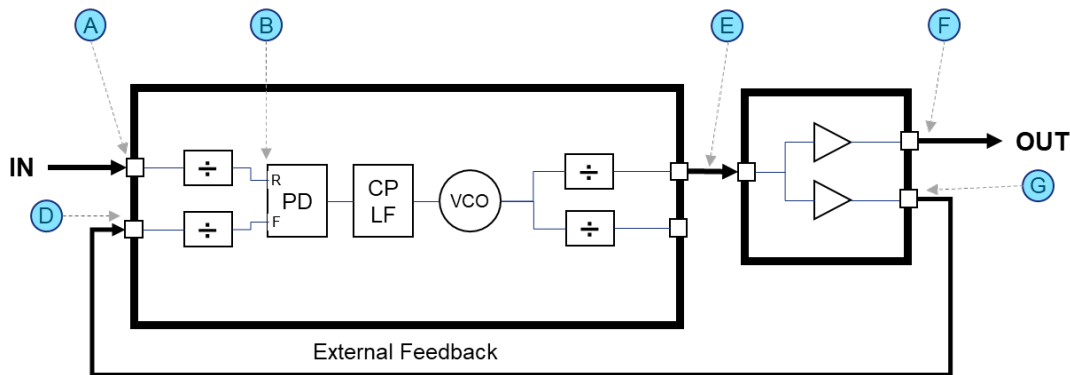
Bringing Fan-out Buffers into the ZDB Loop

In many systems, the outputs of a ZDB PLL will not provide enough copies of a signal for all destinations within the system or the ZDB device might have a large output-output skew specification. In either case, an additional fan-out buffer can be added to generate enough copies or to lower the output-output skews between those copies, since simple fan-out buffers often have much lower output-output skew specifications compared to PLLs.

However, the problem with adding a fan-out buffer after a ZDB is that the propagation delay (t_{PD}) of the buffer delays the signals by a variable amount. Also, the variation between the minimum and maximum t_{PD} adds to the system time error budget. There is also the issue that if multiple fan-out buffers are used then part-part variations between two similar devices can also add to the time error budget.

The solution to these issues is to bring the fan-out buffers inside the ZDB loop(s). If External Loopback is taken from the output of the buffer instead of the output of the ZDB device, then the propagation delay of the buffer will be factored out by the ZDB process. Since the ZDB process is dynamic and on-going, the variation in t_{PD} is also factored-out. Referring to Figure 3, the normal ZDB device would align, Points A, D and E. However, with the fanout buffer added, the External Feedback is tapped at Point G instead, causing Point A, D and G to align. Point F will differ in time from Point G only by the output-output skew of the buffer.

Figure 3. Fan-Out Buffer Inside ZDB Loop



If multiple fan-out buffers are used, each with its own ZDB loopback, then part-part delays are also eliminated, leading to even tighter alignment.

Round-Trip Time (RTT) Measurement and Compensation

The ZDB method is powerful and precise but is not always feasible, especially for signal paths that cross the system backplane. For those cases, a slightly different technique is used.

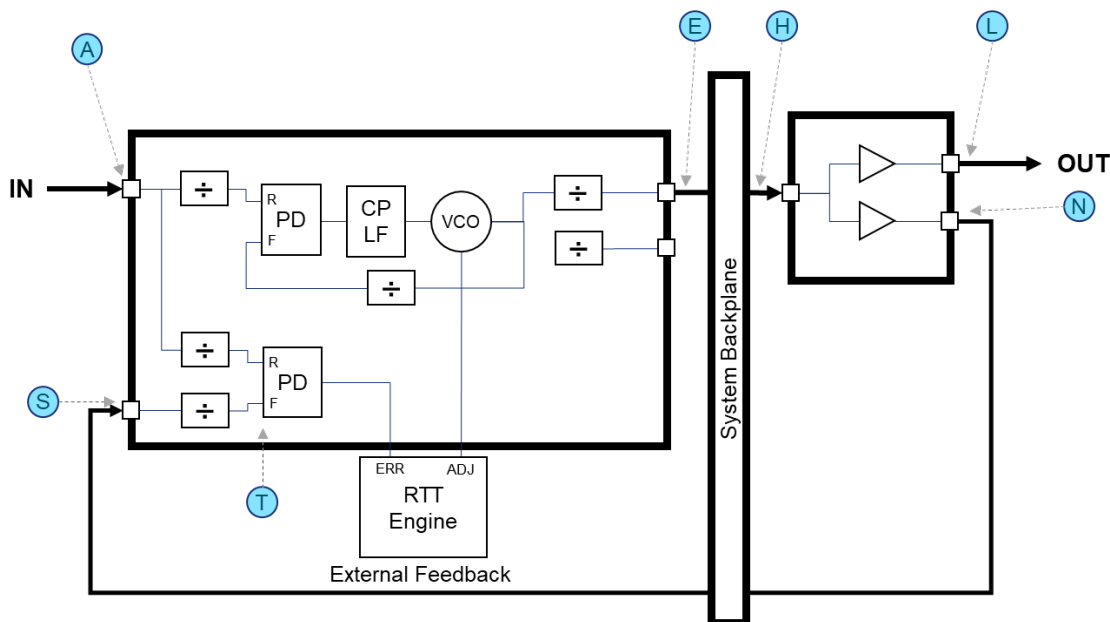
The Round-Trip Time (RTT) measurement and compensation method is very similar to the ZDB concept listed above, specifically for the case where the path delays of the External Loopback versus the forward signal path (IN) are different.

In the ZDB case, shown in Figure 1, alignment was achieved at Point B, the input of the PD, but offsets were inserted into the PD-measured phase error to compensate for the difference in path delays from the inputs, which had the effect of creating alignment at the input pins (Points A and D). Furthermore, a delay of one clock period was inserted in the External Loopback path to create alignment between Points A and C. Finally, a compensation value for the output-output skew was inserted to ensure alignment between Points A and E (IN and OUT).

By knowing the delays in the various paths, compared to our measurement point (the PD), compensation values can be inserted to achieve alignment at any point in the system.

The desire is still to align OUT with IN, but now across a long signal path which can include multiple elements, including a system backplane, buffers and PLL devices. This is shown in Figure 4. Point A is to be aligned with Point E. In the RTT method, the same principle is used, but an extra element is added, called the RTT Engine in the below diagram. The RTT Engine reads one or more error signals (ERR) from Phase Detectors measuring different parts of the signal propagation path (next section discusses how to segment the path for maximum accuracy). It performs the necessary calculations and generates a phase adjustment value (ADJ) that is applied to the forward propagation path of the IN signal. In this example, IN is shown as using a simple PLL path, which can be adjusted as described in IDT AN-1030.

Figure 4. A Simplified Illustration of Round-Trip Time (RTT)



In Figure 4, the RTT is using a phase error (ERR) signal generated by measuring the full round-trip time of the signal and comparing it to the IN signal directly. It adds or subtracts known path delays in the system, such as the difference in delay from A to T versus S to T, the output skew from L to N, etc. This creates the total adjustment value (ADJ) that is applied to the forward path for the IN signal at the VCO, resulting in alignment of A to N across the backplane.

Extending the ZDB and RTT Methods to a Full System

The ZDB and RTT method discussions illustrate the concepts for maintaining tight phase alignment, but in a simple use-case of a single input to a single output. Extending this to a typical chassis-based system without modification would be prohibitively expensive. This section will examine the use of ZDB and RTT concepts together to deal with a system with 1+1 redundant timing cards, where the IEEE 1588 servos and master ToD counters reside and 16 line cards (as an example) that contain PHY devices with embedded ToD counters. This is illustrated in Figure 5. The PHY devices will also recover clocks from the line interface to send to the timing cards for potential selection as the master system clock (SyncE).

Figure 5. Distribution of a Synchronization Pulse Across a Multi-Card System

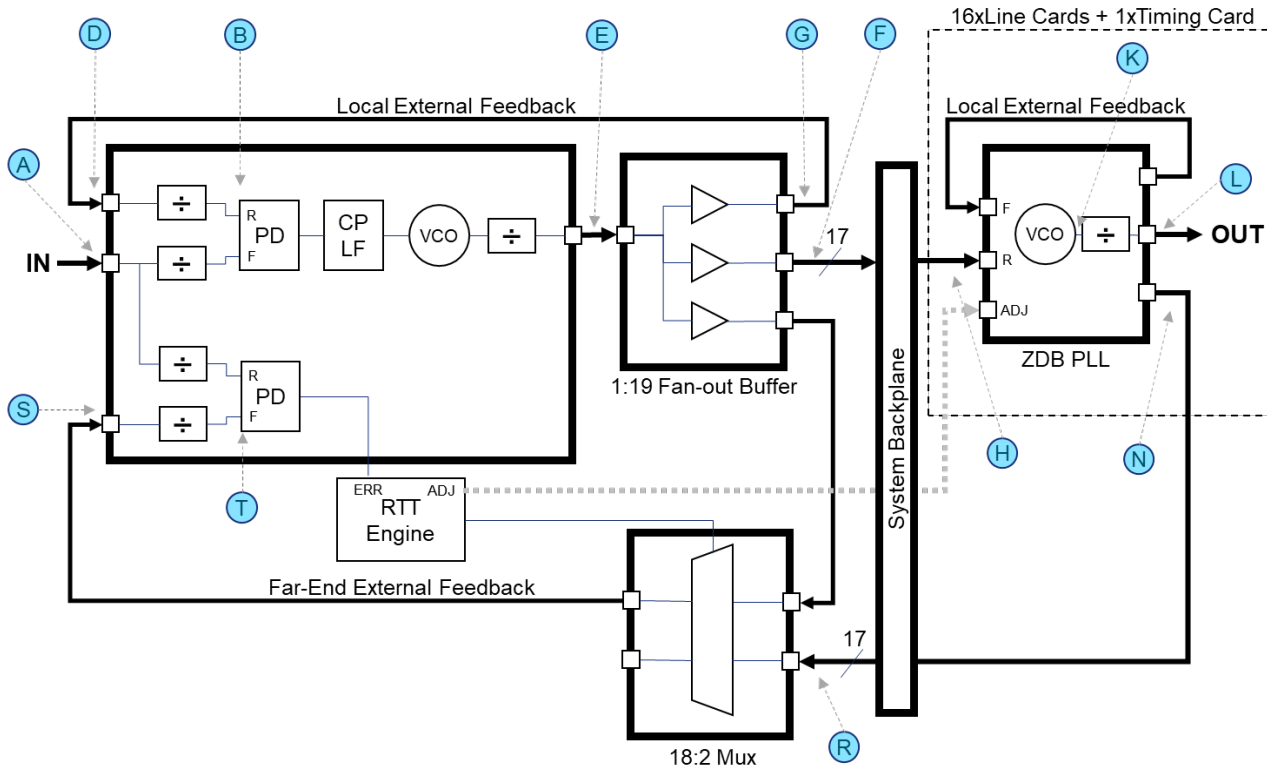


Figure 5 shows the distribution of the Synchronization Pulse across the system.

- The IN signal on the left would be generated by the Master ToD within the SETS function (not shown) on the active Timing Card.
- The OUT signal on the right would be the Synchronization Pulse going to the PHY ToD load signal on the line-card PHY (not shown).
- A copy of the looped-back signal at Point N would also go to the backup Timing Card (not shown).
- The destination is shown using a ZDB to recover the Synchronization Pulse and generate at least one extra copy with minimal uncertainty relative to receipt of the Synchronization Pulse at that card. The destination could also use a ZDB with a fan-out buffer inside the ZDB loop if many copies of OUT are needed on the local card, for example if there are multiple PHYs on the card.
- The illustration also uses a ZDB on the active timing card, but with a fan-out buffer inside the loop to generate the multiple copies of the Synchronization Pulse that are being distributed from the active Timing Card to the destination cards.

What is new in this diagram is the use of a mux on the Far-End External Feedback path. This mux allows the RTT Engine to cycle through each destination card in the system, one at a time in a round-robin fashion. This saves many resources in the system such as multiple PDs and RTTs. Note that the ADJ output of the RTT does not go to the local PLL now. The adjustment values are unique to every different destination card in the system and so need to be applied at the destination, not the source. This is because the source is sending the Synchronization Pulse with identical alignment to each destination.

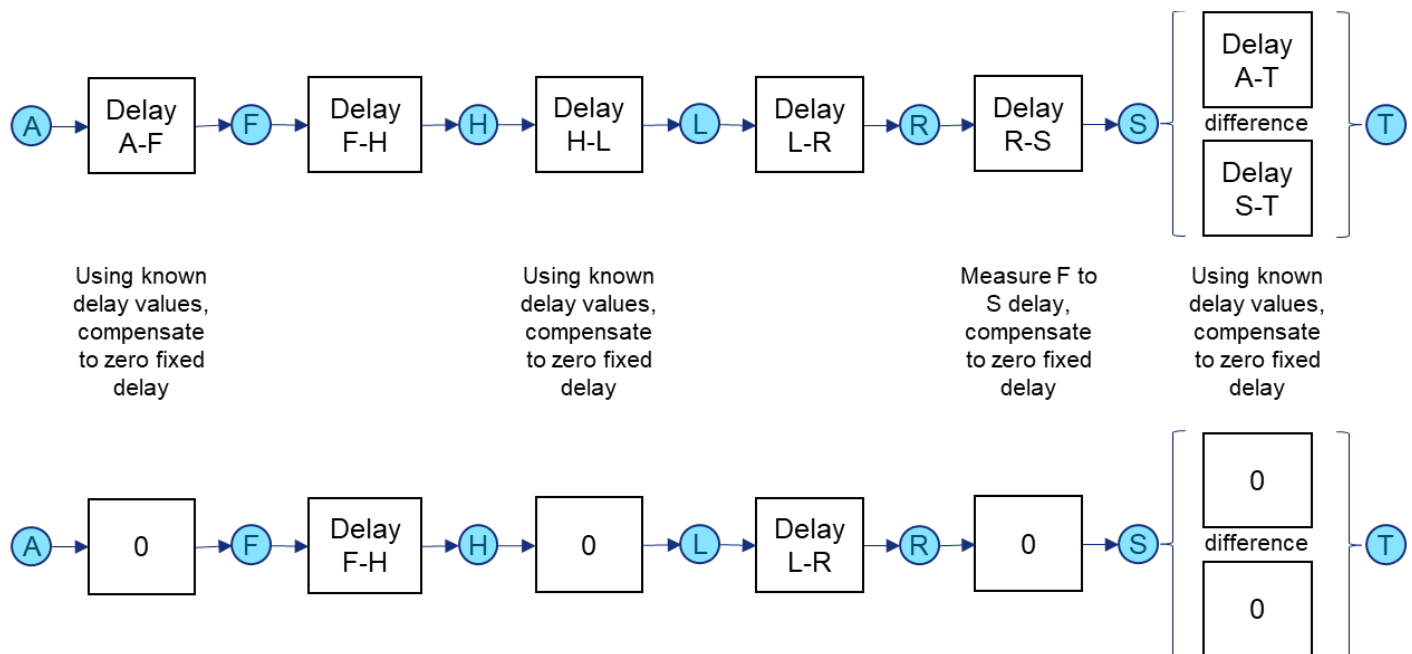
The round-robin cycle must run frequently enough to be able to notice and compensate for significant delay changes as they happen in the system. The most common such disturbance is changes in temperature. Anecdotal information talks about temperature change rates in outdoor equipment (it is usually far less in indoor equipment) of 10-20°C / min. Since a 10°C change is rarely enough to significantly change the system delays, a cycle time of once per several minutes is likely enough.

The ADJ signal from the RTT is shown as a dotted line because this is in the form of a digital word, not a discrete signal and needs to be passed over a data channel to each destination card. Figure 5 is also showing the ADJ signal going into the adjustment input of the ZDB. That is one implementation choice that can be made with respect to the way that delay compensation is done in the system.

As mentioned earlier, one method involves adjusting the position of the Synchronization Pulse to load a fixed ToD value at the right time into the peripheral ToD counters. This is what is shown here. The other method involves keeping the Synchronization Pulse constant and adjusting the ToD value that is loaded into each peripheral ToD counter. If that method were to be used, the ADJ value would be applied by software to the ToD value that was separately transmitted to each destination.

When calculating the ADJ value for each destination, the RTT Engine is attempting to calculate the delay from Point F to Point H in the system and compensate for the delay present at that point in time. As discussed previously, the ZDB fixed delays can be compensated for and zeroed-out. This is shown in the progression from top to bottom in Figure 6. Delay values for specific components should be provided by their suppliers in a datasheet or application note, such as AN-1035 for IDT's ClockMatrix devices.

Figure 6. RTT Engine Calculation Overview



To compensate for the delay from Point R to Point S, the additional trace shown on the active timing card, from the fan-out buffer at Point F to the mux input at Point R can be used. Delay F-H (forward path across the backplane) will be nominally the same as Delay L-R (feedback path across backplane) if PCB trace lengths from E to H are similar to the backplane connection to Point H. The mux input-input skews should also be considered in this calculation. However, with some care and known data the fixed delays in the full round-trip can be reduced to just Delay F-H and Delay L-R. These are the outbound signal path and the return signal path respectively. Since these consist of just PCB trace delays, it is possible by examining the layouts to know the proportional relationship between the outbound and return signal paths. Most systems would assume equal path lengths, but if it is known that the relationship is, for example, 60 / 40%, then that relationship could be used instead. Using that information, the RTT Engine can use the ERR input value measured at Point S to determine Delay F-H. This is then communicated as the ADJ value for that destination card. Applying that adjustment in the destination, either to the Synchronization Pulse or to the peripheral ToD value, the master and peripheral ToDs can now be aligned on all cards with a high degree of precision.

For the RTT calculation to work for each line card in the system (or from master to the slave timing card), the round trip time must be a deterministic number of Synchronization Pulse clock periods. For example, the path to any specific line card the round trip time must be always less than one period or always greater than one and always less than two periods, etc. If the delay can span from one period into a later period over some range of conditions, the Synchronization Pulse clock rate must be slowed down to ensure the relationship becomes deterministic. Otherwise it would be impossible to calculate the ADJ value properly. Note that if each is deterministic, they do not need to be the same for each line card or timing card path.

As one more potential refinement, if a second PCB trace is added on the timing card from the fan-out buffer to the mux, but with a different PCB trace length, then the RTT can measure the two paths and calculate a delay value for PCB traces per unit length at that time / condition. This can be used to compensate for trace length mis-matches or for changes induced by different environmental conditions for any PCB traces on the boards.

Summary

This application note has shown how to use ZDB and RTT techniques together in a multi-card system to achieve tight alignment between a master and several peripheral ToD counters across backplanes and fanout buffers with highly variable delays. Devices such as IDT's ClockMatrix family, with the necessary precision phase adjustment, zero-delay buffer capability and precision phase measurement capabilities are needed to achieve some of the tighter ITU specified alignment classes.

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 may 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.