Order Number: AN1934/D
Rev 0, 09/2001

AN1934

Effects of Skew and Jitter on Clock Tree Design

Prepared by: Don Aldridge and Tom Borr
September, 2001

ABSTRACT

This application note discusses the parametrics of skew and jitter as these terms apply to PLL clock drivers and clock buffers. The application note covers the definition of the various types of skew and jitter, the measurement techniques and values associated with these parameters, and concludes with an example clock tree design and analysis of the skew and jitter.

INTRODUCTION

At first glance, clock distribution trees are relatively simple. As shown in Figure 1, a typical clock distribution tree consists of a clock source and a series of clock distribution buffers that deliver multiple copies of the clock source to many locations in an electronic system. The clock source may be a crystal oscillator (Figure 1) or an external clock source. This clock source may be at the desired frequency or may need to be translated to the desired frequency or frequencies as part of the clock tree circuitry. The clock tree will consist of some combination of PLL clock drivers and/or fanout buffers providing multiple outputs. The clock tree may consist of several devices or be composed of a single integrated circuit. Individual clock outputs deliver the clock signal to various locations on a PC board.

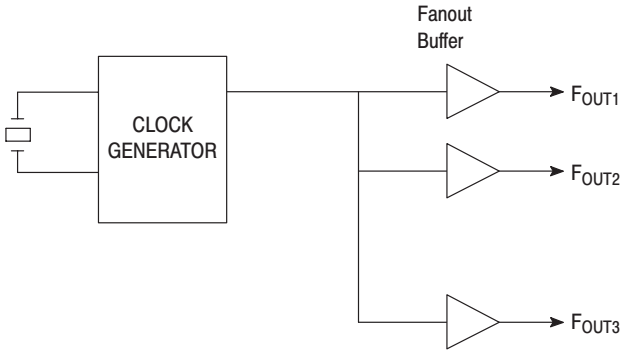

Figure 1. Typical Clock Distribution Tree

After a more detailed look at the requirements of the clock system and the data sheet specifications of the devices used to implement the clock tree, the design appears a bit more complicated and requires a more detailed analysis. The specifications that are most important for this type of analysis are usually found in the AC parameter portion of the clock driver data sheet and consist of parameters such as propagation delay, skew, and jitter.

This application note discusses these parameters by reviewing the definition, the measurement techniques, and their effects on system performance. The application note concludes with the analysis of a typical clock distribution tree based upon these parameters.

Standards for Skew, Jitter Definitions, and Notations

The reference used to determine the standards for skew, jitter definitions, and notations for this application note are the EIA specification EIA/JESD65. This EIA specification documents the current industry standard for these parameters. This document is available in a downloadable PDF format at the web site of http://www.jedec.org/.

Clock Driver Devices

Clock driver devices consist of both clock fanout buffers and PLL based clock generators. Typical fanout buffers are shown in Figure 2 and consist of an input buffer driving many outputs though individual output buffers. The specific devices shown are the MPC942C and MPC942P which offer 18 LVCMOS outputs and have either a LVCMOS input or a LVPECL input. Some fanout buffers have an optional internal divider network to produce an input clock divided by two. A Phase Lock Loop device is shown in Figure 4 and may or may not have built–in fanout buffers. It is not the intention of this application note to explain all of the fundamentals of a PLL; however, a quick review of the components is covered.

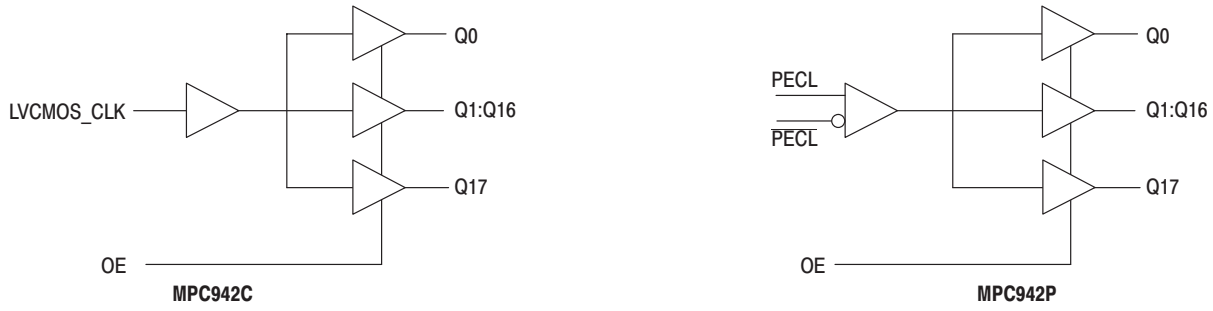

Figure 2. Fanout Buffer

A basic PLL clock architecture (Figure 3) consists of a phase detector, a low pass filter, a $V_{CO,}$ and (in this diagram) two divider networks. Both dividers are at the inputs of the phase detector. The input to the clock driver, or the reference frequency, may be external or sourced from a crystal oscillator that is included as part of the clock driver architecture. This input frequency may be divided by an optional P divider block and then applied to the input of the phase detector. The phase detector produces a correction signal based upon the difference in phase in its two inputs. The correction signal or the output of the phase detector is filtered and applied to the input of the voltage controlled oscillator; V_{CO}. The output of the V_{CO} is applied to the M divider and becomes feedback and the second of the two inputs to the phase detector. When the loop is in "lock," the two inputs to the phase detector are the same frequency and the same phase. The output frequency, or $F_{OUT,}$ is the reference frequency divided by P and then multiplied by M and will continually track the reference frequency.

With a few additions to the basic PLL clock architecture, we can create a multi–frequency and multi–clock distribution device as shown in Figure 4. The more complex divider network shown provides the M divide value for the feedback path to the input of the phase detector and, also, the N divider divides down the V_{CO} frequency to the desired system frequency or frequencies. Multiple outputs from the clock divider may provide for the generation of multiple frequencies. Note that fanout buffers are included for each output to provide the required system clock drive. Also note that an equivalent fanout buffer is included for the feedback path. The feedback connection for the PLL is external to the device and thus equalizes the delay through the main clock outputs. As is discussed later, the external feedback path may also include compensating trace delay which allows the phase of F_{OUT} clock to be advanced forward or backward with respect to the input clock.

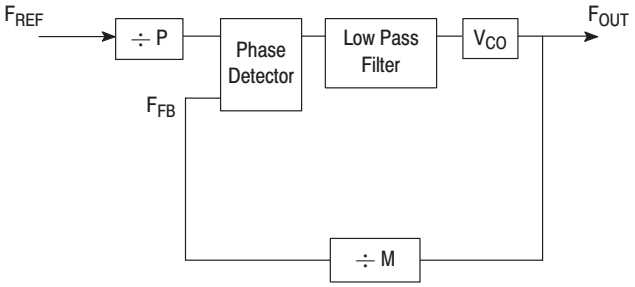

$F_{OUT} = (F_{REF} / P) * M$

Figure 3. Basic PLL

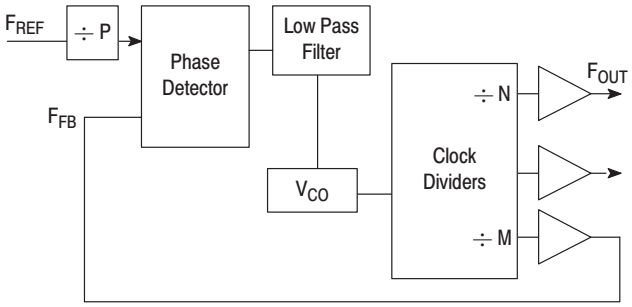

$F_{OUT} = (F_{REF} / P) * (M/N)$

Figure 4. PLL Based Clock Driver

Clock Driver Parametrics

The clock driver parameters that are of interest in this application note are Buffer Propagation Delay, Zero Reference Delay, Skew, Jitter, and PLL Bandwidth and Jitter. These parameters are typically found in the AC parameter portion of a clock driver data sheet.

Fanout buffers have output skew, jitter, and propagation delay. PLL clock driver devices are characterized with jitter, output skew, and an effective input to output propagation delay called Zero Reference Phase Delay. In a clock tree design, the parameters that complicate the analysis is skew between the outputs of a clock fanout buffer and edge or frequency jitter. Jitter commonly is generated in the very early stages of a clock tree and potentially at each stage of the clock tree. These jitter sources may or may not be cumulative and be passed to the outputs.

Buffer Propagation Delay

Clock distribution buffers have a propagation delay from input to output. Typical values for this delay are in the order of a few nanoseconds. Data sheet specifications may be given for a single propagation delay or as separate values given for a low to high edge; versus a high to low edge. The low to high edge value and the high to low edge value should be very similar but not necessarily the same. The notation for propagation delay is t_{pd}. Also the notations of t_{plh} and t_{phl} are used to indicate the propagation delay for a low to high transition and a high to low transition of a waveform, respectively.

Reference Zero Delay

Reference Zero Delay is the JEDEC term for the effective PLL buffer delay. This parameter is also referred to as Static Phase Offset (or SPO). The JEDEC notation is $t_{(\phi)}$. The value of SPO is defined as the average difference in phase between the input reference clock and the feedback input signal, when the PLL is locked. The value of the Reference Zero Delay can be compensated for by including PC board trace delay in the feedback path of the PLL. Specially constructed PLL clock drivers called Zero–delay Buffers make use of this occurance and can produce a clock edge at the output that is exactly in phase with the input.

Skew

Clock fanout buffers and PLL clock drivers with built–in fanout buffers offer multiple outputs. These clock outputs are routed across a PC board to various devices. A typical fanout buffer may have as many as 18 to 20 clock outputs. Typically, these outputs are designed to drive a 50 ohm cable or a 50 ohm PC board trace. Ideally, all of the outputs are timed such that clock edges on each output switch at exactly the same time. However, real life devices do not. Small amounts of skew exist between the high to low or low to high transition on one output as compared to another output. For systems that require synchronization between data and clocks or multiple clocks on the PC board, this skew is a bad thing. Clock integrated circuit designers try to minimize the amount of skew in a device. However, skew does exist and the device data sheet usually specifies the amount of skew.

This output skew is typically defined in three ways: output–to–output, process, or part–to–part skew.

Output–to–output skew is defined as the skew between the various output edges on a single device. Process skew is defined as the skew between the same output pin on two different devices. Finally, the part–to–part skew is defined as the skew between any output on two different devices. Figure 5 illustrates output skew types for both single–ended and differential output waveforms. Typically, both output–to–output and part–to–part skew are specified on a data sheet. The JEDEC specification states that the skew values are to be determined with the outputs driving identical specified loads.

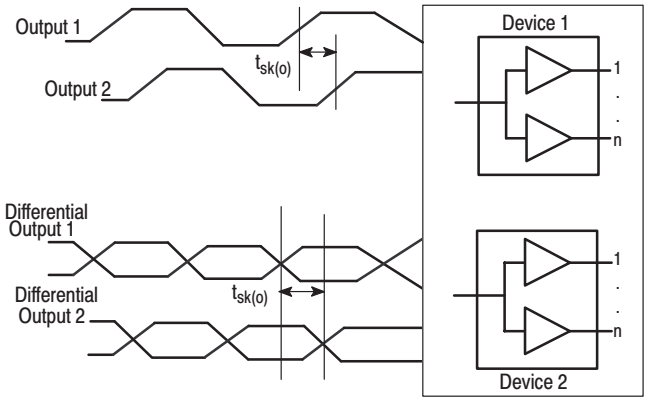

Figure 5. Output–to–Output Skew

Jitter

Jitter is a deviation of the edge location on the output of the clock buffer. As with skew, jitter is a bad thing and is usually measured in picoseconds. There are three categories of jitter that are of interest: cycle–to–cycle, period, and phase jitter.

Cycle–to–cycle jitter is the difference in the period of any two adjacent clock cycles. The difference is reported as an absolute value according to the JEDEC specification. However, quite often a ± value is used. The JEDEC symbol for cycle–to–cycle jitter is $t_{jit(cc)}$. Cycle–to–cycle is usually measured over some large sampling of cycles and specified as the maximum difference. Figure 6 shows the measurement and calculation of cycle–to–cycle jitter.

In PLL based systems, the value of the cycle–to–cycle jitter is usually small since the PLL does not quickly respond to changes on its input. Since cycle–to–cycle jitter is the difference in the period from one cycle to the next, this jitter is at the clock frequency. This jitter is also referred to as short term jitter.

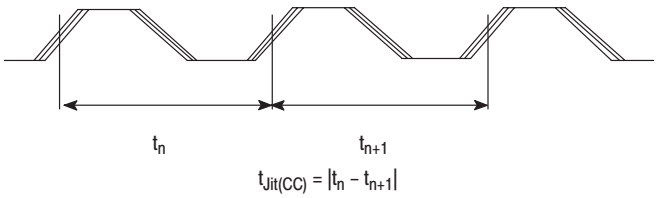

Figure 6. Cycle–to–Cycle Jitter

Clock integrated circuits have an inherent jitter generated within the device. In addition, external sources contribute to this jitter. Specifically, power supply noise may be a source of jitter in both PLL based and non–PLL based clock driver devices. Power supply design, power supply filtering, and board layout contributes to the overall jitter values measured on the output of the clock device.

The second type of jitter is period jitter, which is defined in the JEDEC specification as the deviation in cycle time of a signal with respect to an ideal period. Figure 7 shows the definition and calculation of period jitter. This jitter type is reported as an absolute maximum value as measured over a long time period. The JEDEC symbol for period jitter is $t_{jit(per)}$. The long time period varies from measurement system to measurement system. Typical time periods are 64 microseconds which, at a frequency of 100 MHz or so, yields many (6400) clock cycle period values. This type of jitter represents the random movement in the instantaneous output frequency or output period of the clock source.

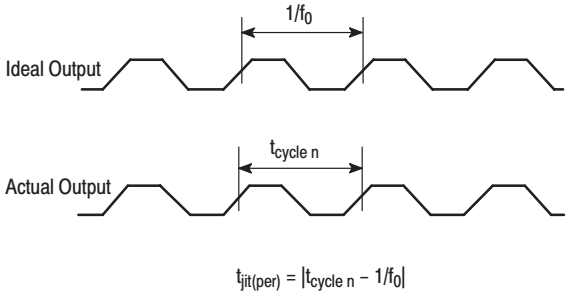

Figure 7. Period Jitter

The last jitter type covered is that of phase jitter. Phase jitter is associated with PLL based clock drivers. The JEDEC specification notation is $t_{jit(\phi)}$. This value represents the input to output jitter associated with a PLL clock driver. The value is given as an absolute value of the range or variation in the difference between the phase of the reference input and the phase of the feedback input to the integrated circuit. (See Figure 8)

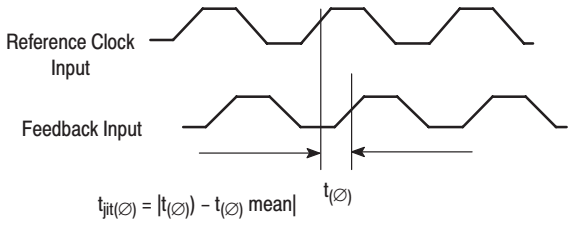

Figure 8. Phase Jitter

Jitter Values and Data Sheets

Definitions of the various types of jitter and the equations for calculating the values are necessary for an understanding of jitter and how it relates to a clock design. However, another important factor in understanding clock drivers is the metrics and methods used of the published specifications. These metrics may vary from manufacturer to manufacturer and also within a manufacturer's clock driver offering. The values may be given as RMS values or as a peak–to–peak value. They may be listed as typicals or as actual maximum values. Understanding of the background of the jitter values on a data sheet are necessary for circuit design as well as comparing clock driver devices.

Jitter measurements, whether cycle–to–cycle period or phase, are measured over some large number of samples. The data for a typical device, when plotted, represents a classic distribution Gausian or bell shaped curve where most of the clock cycles are close to the ideal frequency (in the case of period jitter) with fewer and fewer devices having increasing deviation from the ideal period.

In classical statistics, the distance from the center of the Gausian curve is defined in values of standard deviation or sigma; and the higher the sigma multiplier, the higher the confidence level is that a device will not exhibit a jitter value greater than a predefined amount.

Data sheet specifications that list RMS values imply a 1 sigma deviation above the mean and 1 sigma deviation below the mean or a total of 2 sigma confidence level. Table 1 lists these confidence factors for ±1 sigma through ±6 sigma. If data sheet values are specified as RMS values and higher levels of confidence are desired, then the data sheet values for jitter must be multiplied by the desired confidence factor.

Figure 9 shows the Gausian distribution curve and sigma points for a device with an output frequency of 400 MHz (period of 2.5 ns). A value of ±3 sigma or 6 sigma gives a confidence level or a probability that the clock edge is within the distribution of 0.9970007%. The ±3 sigma limits define the upper period limit of approximately 2.52 ns and the lower period limit of approximately 2.48 ns.

Figure 10 shows actual measurements made for the period jitter of a 400 MHz clock device.

In this example, the mean period is 2.49921 ns (approximately 2.5 ns) with the standard deviation being 6.48 ps and the peak–to–peak jitter being 57 ps. This data was captured with a sample size of 13100 samples.

Table 1. Confidence Factor

Sigma	Value	Confidence Factor
±1	(2 sigma)	0.68268948
±2	(3 sigma)	0.95449988
±3	(6 sigma)	0.99730007
±4	(8 sigma)	0.99993663
±5	(10 sigma)	0.99999943
±6	(12 sigma)	0.99999999

RMS values for jitter look better on the data sheet than peak–to–peak. However, peak–to–peak values may be needed for clock jitter analysis. If RMS values are specified, peak–to–peak values may be derived based upon the required system reliability for the specific applications. If the other cases where peak values are specified, these values may be used directly. However, the question that must be asked in order to use the manufacturer's peak–to–peak values is what level of uncertainty is being specified. These parameters may be specified differently on each manufacturer data sheet. Therefore, care must be taken to insure that when one is comparing values, the values are specified and measured in a similar fashion.

The above discussion assumes the jitter measurements have the classic Gausian curve or statistical distribution. This would be the case if the jitter is completely random. However, clock driver devices may have internal mechanisms that produce jitter that deviates from the classic bell curve. In this case, the total jitter is composed of the addition of a series of bell curves providing a more complex distribution. With care, the terms of RMS and the various sigma levels may still apply.

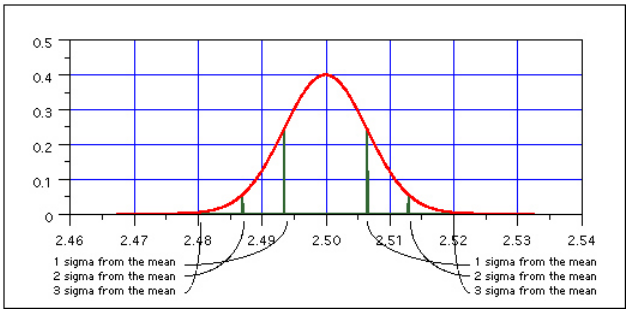

Figure 9. Classic Gausian Distribution Curve

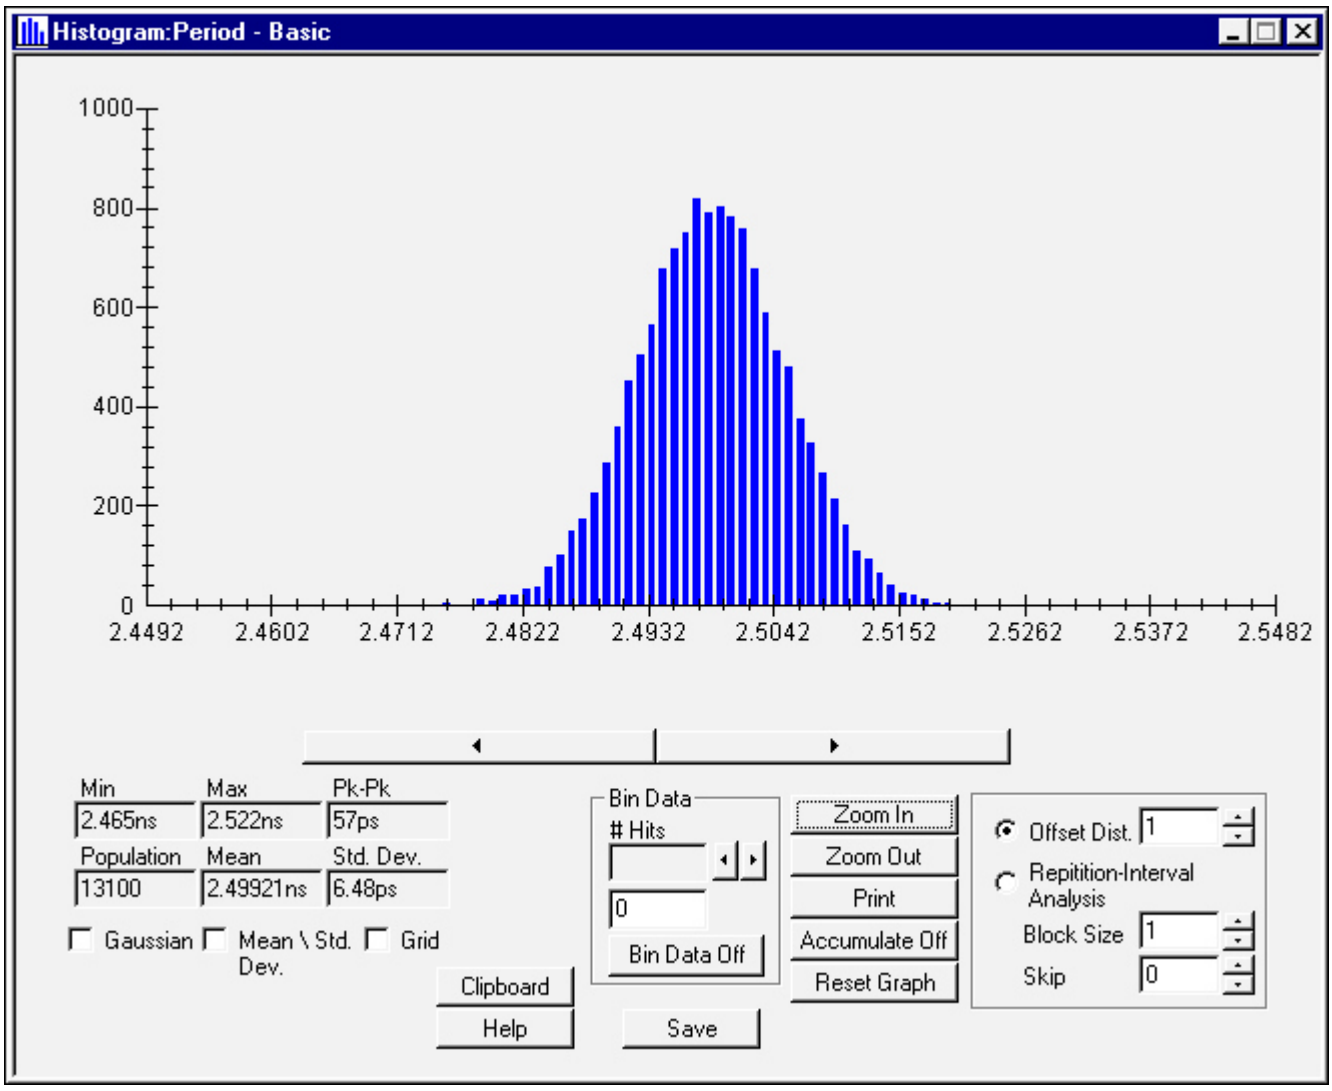

Figure 10. Period Jitter for Typical PLL Clock Device with F_{OUT} = 400 MHz

PLL Bandwidth and Jitter

The PLL based clock driver locks on to a reference frequency and maintains an output frequency based upon that reference frequency. If the reference frequency changes, the PLL attempts to follow the change in the reference frequency. However, if the change is faster than the PLL can follow, the PLL based clock driver acts as a low pass filter and ignores or effectively filters out the higher frequency changes on its input. As with any low pass filters, the PLL has a cutoff frequency, or bandwidth, associated with it. This bandwidth becomes important to our clock tree design. High frequency noise and jitter will not pass through the PLL.

The actual bandwidth of the locked PLL is dependent on many factors, including the feedback divider ratio. The higher the divide ratio, the lower the bandwidth. Thus, those PLL clock driver devices that have selectable feedback divide ratios will have varying bandwidth values. Bandwidth may or may not be specified by the PLL clock driver manufacturer. If not specified, the information is usually available on request.

Figure 11 is a typical PLL frequency modulation bandwidth waveform. Note the cutoff frequency is about 300 kHz. As mentioned before, the cutoff frequency will vary with a change in the divide ratio in the feedback loop. Bandwidths of PLL–based clock drivers vary from low values of a few kHz to higher values of a MHz or more depending upon the intended application of the device. Clock synthesizers typically have the lowest bandwidth. Bandwidths of these devices are in the order of 30 to 50 kHz. Clock generators are next with bandwidths of a few hundred kHz. The devices with the highest bandwidth are the Zero–Delay–Buffers. The bandwidths of these devices are typically a MHz and above.

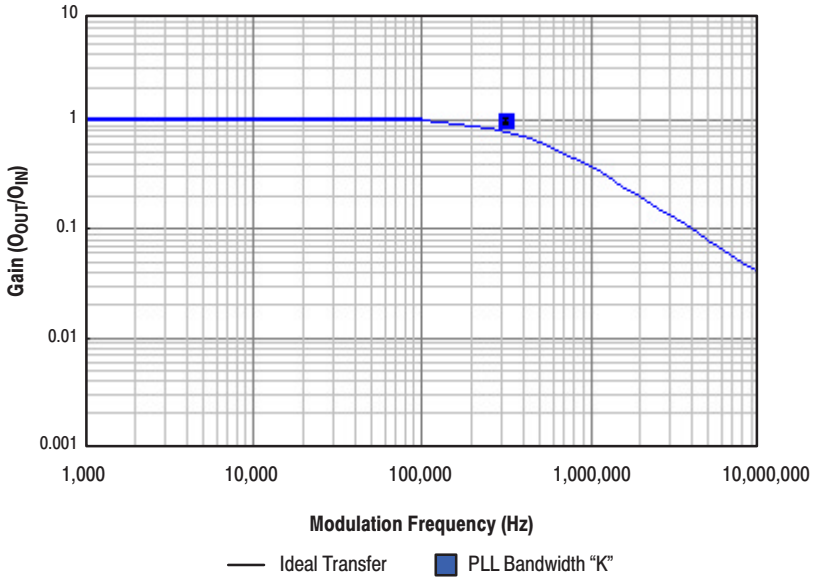

Figure 11. PLL Bandwidth

Clock Tree Application Example Analysis

Figure 12 shows an example clock tree application which is used to show the effects of skew and jitter in a system. The goal of this design is to provide multiple phase aligned HSTL and LVCMOS outputs. The reference for this clock tree is a crystal oscillator while the outputs of the tree provide both HSTL and LVCMOS to various system locations. The tree starts with a crystal oscillator in the MC12430 integrated circuit. The MC12430 is a PLL based clock synthesizer that allows very fine control of the output frequency in 1 MHz steps.

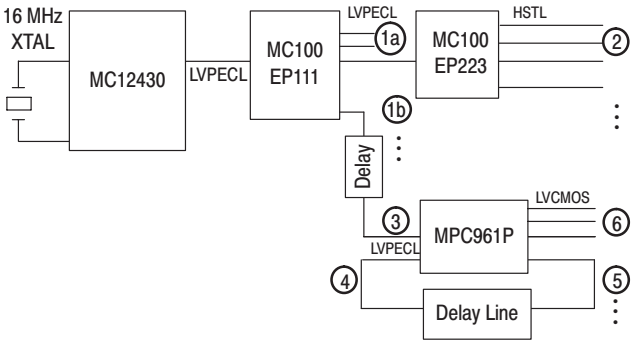

Figure 12. Example Clock Tree

The LVPECL output of the MC12430 drives the LVPECL input to the MC100EP111 fanout buffer.

The MC100EP111 consists of 10 LVPECL differential pairs of which one pair is connected to a MC100EP223 input and another pair connected to the MPC961P input. The MC100EP223 fanout buffer provides HSTL outputs while the MPC961P is a PLL–based clock generator that provides several LVCMOS outputs. Notice there is an introduced delay from the LVPECL output of the EP111 to the LVPECL input of the MPC961P zero–delay buffer. This introduced delay is due to backplane or cable distance which will skew the LVCMOS outputs to later than the HSTL outputs. To compensate for this delay, the MPC961P zero–delay buffer is used in conjunction with a PC board delay line in the feedback path.

Figure 13 provides an analysis of the example clock tree for the situation where we have no (or choose to ignore both) jitter and skew in the devices. Later, in Figure 14, an analysis with both jitter and skew is shown. The circled numbers in Figure 12 are used as reference points on the clock analysis waveforms.

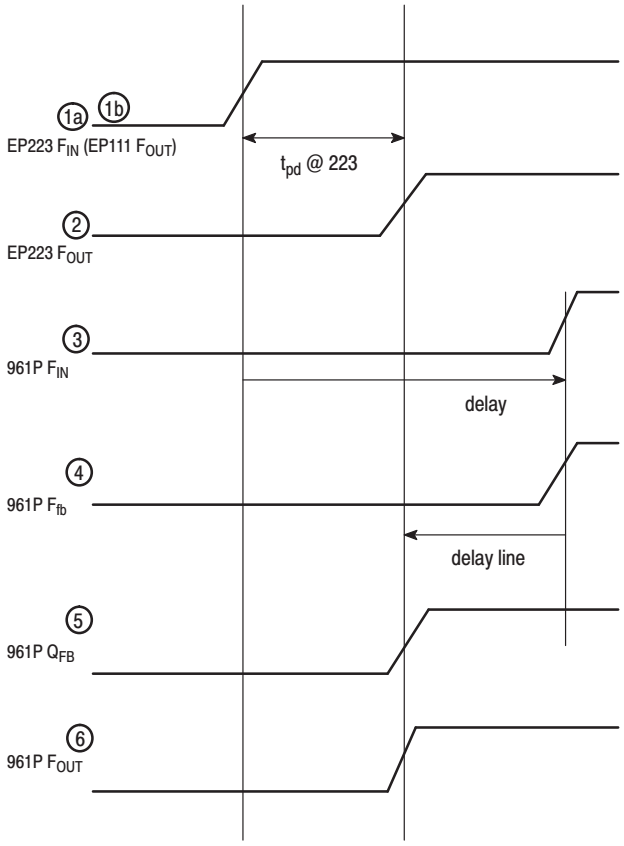

Figure 13. Example Clock Tree Analysis without Jitter and Skew

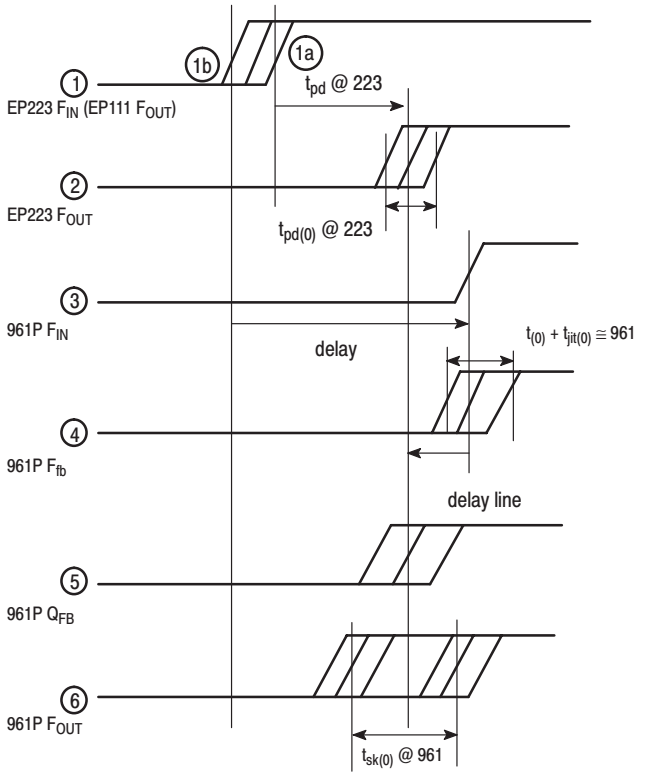

Figure 14. Example Clock Tree Analysis with Jitter and Skew

The first analysis of the clock tree is with the assumption that there is zero output–to–output skew and zero jitter. With zero skew between the outputs of the MC100EP111, the signals 1a and 1b are identical. The waveform at point 2 is delayed due to propagation delay, t_{pd}, of the MC100EP223.

The waveform at point 3 is delayed due to the delay associated with the backplane or cable distribution. By using the MPC961P zero–delay buffer and placing the appropriate trace delay in the feedback path of the PLL, we can compensate for the backplane trace and bring the waveform for points 5 and 6 back in line with the output of the MC100EP223.

Next, we will do the same analysis but we will include the output skew on the MC100EP111, the MC100EP223, and the phase jitter for the MPC961. (See Figure 14.)

Initially, we have the output–to–output skew for the EP111. We will assume that the output connected to the MC100EP223 is the slowest and thus the longest delay output and that the output connected through the delay line to the input of the MPC961C is the fastest or shortest delay output. This analysis must also be done for the opposite situation; where the MPC961P is connected to the slowest output and the MC100EP223 is connected to the fastest output.

The waveform at point 2 is delayed from point 1a by the propagation delay, t_{pd}, of the MC100EP223 as we had in Figure 13. However, by including the output–to–output skew for the MC100EP223, we find uncertainty in the location of point 2 as shown in the waveform of point 2 of Figure 14.

Point 3 shows the waveform arriving at the input to the MPC961P. Point 4 is the feedback input to the MPC961P and, in the ideal case, is exactly the same as point 3. However, in this situation, we have the uncertainty caused by the combination of the static phase offset, $t_{(\phi)}$ and the phase jitter for the MPC961P. Figure 14 depicts these two values added together. The waveform of point 5 now moves back in time due to the delay line in the feedback path of the PLL such that the nominal output now coincides with the HSTL outputs of point 2. Point 6 shows the added uncertainty of the outputs due to the

output–to–output skew from the MPC961P. Note that Figure 14 is not to scale and the magnitudes of the skew and jitter are actually much smaller than indicated.

Table 2 lists the skew and jitter values for the 3 devices used in the example. The MC100EP111 has a specification for output–to–output skew. The MPC961P has parametric values for Static Phase Offset, phase jitter and output–to–output skew. The MC100EP223 has values for propagation delay and output–to–output skew.

Table 2. Confidence Factor

Device	Parameter
MC100EP111	tsk(o): 70 ps
MPC961P	$t_{(\varnothing)}$: –50 to 225 ps $t_{jit(\varnothing)}$ ÂF ~ 100 ps $t_{sk(o)}$ ÂFÂ @ 150 ps
MC100EP223	t_{pd}: estimated 700 ps $t_{sk(o)}$: estimated 50 ps

A similar analysis can be done for the case where the MC100EP111 outputs connected to the MC100EP223 and MPC961P are reversed in time. Thus, the MC100EP111 output connected to the MC100EP223 is the fastest and the output connected to the MPC961P is the slowest.

One last comment on the example is that the jitter associated with the clock source, MC12429, was not mentioned in the analysis. This value was ignored for this example due to the fact that all of the clock outputs are derived from the same source, and jitter that occurs on one output would show up on all outputs. Applications that have clock outputs derived from different sources, or have the clock source compared to a frequency standard, would mandate the need to include the source jitter in the analysis.

SUMMARY

Skew and jitter are real characteristics of clock driver devices and may or may not be of importance to a clock tree design. Understanding data sheet values and applying this knowledge to clock tree design can sometimes be a real challenge. With a bit of analysis, one can determine which parameters are critical to a specific clock tree design and be able to compare the values of these parameters from one device/vendor to another. Lastly, data sheet values of jitter are often measured in a lab environment under the best of conditions. Real designs with clock drivers on PC boards with other digital circuitry, noisy power supplies, long traces, and other clock sources can make the overall jitter worse. Careful design practices are a must for the best clock driver design.

NOTES

NOTES

NOTES

How to reach us:
USA/EUROPE/Locations Not Listed: Motorola Literature Distribution; P.O. Box 5405, Denver, Colorado 80217. 1–303–675–2140 or 1–800–441–2447

JAPAN: Motorola Japan Ltd.; SPS, Technical Information Center, 3–20–1, Minami–Azabu. Minato–ku, Tokyo 106–8573 Japan. 81–3–3440–3569

ASIA/PACIFIC: Motorola Semiconductors H.K. Ltd.; Silicon Harbour Centre, 2 Dai King Street, Tai Po Industrial Estate, Tai Po, N.T., Hong Kong. 852–26668334

Technical Information Center: 1–800–521–6274

HOME PAGE: http://www.motorola.com/semiconductors/