Spread Aware™

The use of Spread Spectrum Timing (SST) technology has pervaded the motherboard market to the point where it is being used in virtually all designs using chipsets that support a 100-MHz Front Side Bus (FSB). Spread spectrum timing signals are used for PCI, CPU, and memory buses. When AGP signals are provided by the main system Frequency Timing Generator (FTG), they are also spread. All motherboard chipset vendors are designing their parts to work with spread spectrum timing signals. While the fast-paced motherboard market has quickly adopted the technology, it has also been permeating other markets for quite a while and is promptly being embraced.

The technology was developed solely for the purpose of reducing peak EMI. It eases passage of regulatory testing, and reduces total manufacturing costs. Historically, EMI problems have been dealt with after a board has been designed and its functionality proven. If prototype boards failed EMI tests, EMC engineers would step in and fiddle with component values, add some combination of chokes, more shielding or better gaskets to address specific frequencies at which the board failed.

Prior to the advent of SST technology, the only element a designer could use to help this effort would be to add power and ground layers to the board, a relatively expensive choice and one requiring a redesign. Designers who have utilized SST technology have typically found no less than 8 dB of reduction in peak EMI, and as much as 20 dB on higher order harmonics. Also, while most traditional methods of addressing EMI at the circuit level only benefit a specific harmonic, SST technology is applied to the fundamental frequency and affects every harmonic. The amount of benefit varies from design to design, and there are several factors that influence it.

Figure 1 shows a block diagram of the concept, not an actual implementation. Although the method of implementation varies from supplier to supplier, the concept remains the same. Basically the frequency of the output is slowly swept through a range of frequencies rather than being delivered at a single constant frequency. Two factors (modulation width and modulation profile) significantly affect the amount of reduction in peak EMI that the system will achieve through the addition of SST technology. The wider the modulation, the larger the band of frequencies over which the energy is distributed and therefore the more reduction from the peak you will see. The effect of the modulation profile is less obvious, but in some instances is equally important.

Three modulation profiles have been explored, and two are commonly used in Spread Spectrum FTG products. *Figure 2* shows the typical spectrum analyzer displays of the modulated output signals resulting from the three modulation profiles. *Figure 2(c)* is the resultant spectrum when using a sine waveform for modulation; although easiest to implement, the extreme peaking at the edges reduces the benefit so much that nobody is using this profile in their products. *Figure 2(b)* shows the spectrum when using a triangular (a.k.a. linear) modulation waveform; again, you see a slight trough in the middle of the profile with peaking at the edges. Some FTG manufacturers have implemented this profile. *Figure 2(a)* shows a flat profile, which is optimal, resulting from use of the modulation waveform shown in *Figure 3*. Various terms have been used to describe this wave shape (optimized, Hershey's Kiss, non-linear, et al.).

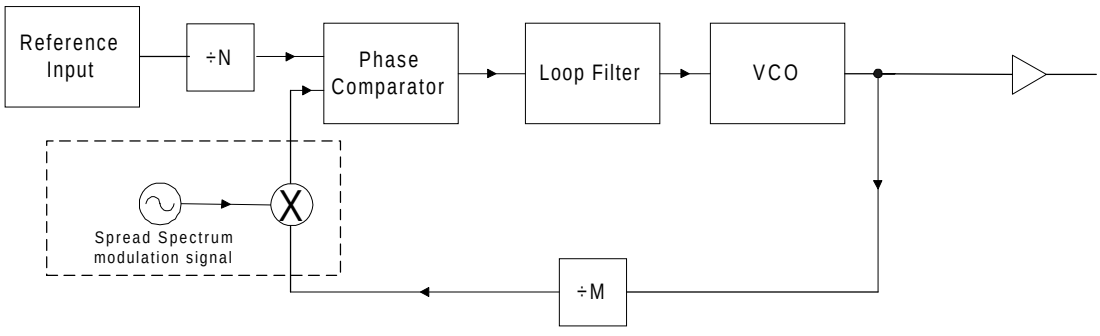

Figure 1. Simplified Block Diagram of Spread Spectrum Timing Technology

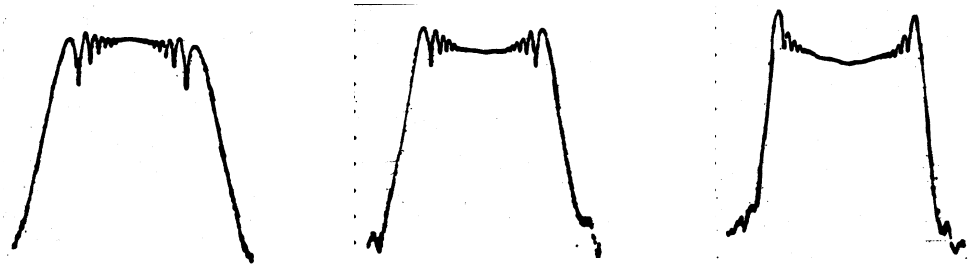

(a) Optimized Modulation (b) Linear Modulation (c) Sinewave Modulation

Figure 2. Frequency Spectra of Timing Signals Resulting from Different Modulation Profiles

Why is a flat spectral profile optimal? The band of frequencies over which the EMI energy is spread is fixed by the modulation width used, it does not vary with the modulation profile. Since EMI testing is concerned with peaks, you want to spread that energy evenly over this frequency band. A flat spectral profile with minimal peaking shows that the energy is evenly spread across the frequency band.

There are a variety of products on the market using the optimal profile. Some, like the Cypress W42C31-03 or W181, have small footprints and can easily be placed in an existing timing path to add the spread spectrum function to it. Others, like the Cypress W144, use a 14.318-MHz reference crystal and provide all the timing signals required to drive the CPUs, core logic, SDRAMs, USB and ISA interfaces. All of these devices use the optimized profile so they provide maximum peak reduction over the entire frequency band.

Spread spectrum timing is a very effective tool for reducing peak EMI, and may be easily integrated into many different systems without affecting other circuit elements. The one type of circuit element that may cause a timing problem when driven by a spread spectrum timing signal is a downstream phase locked loop (PLL).

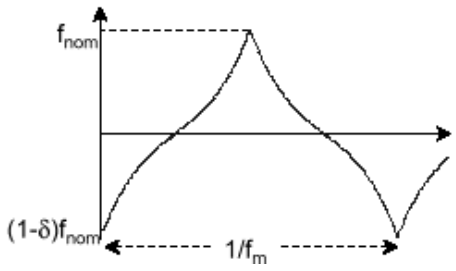

Figure 3. Optimized Modulation Profile

A downstream PLL is a device that receives a reference timing signal from another PLL-based device, including those that utilize SST technology. The term downstream may also indicate that data signals have been transmitted by a processor and the PLL must accurately recover the timing information in order to correctly latch the data. For this reason tracking skew is very important in downstream PLL applications. Some examples of downstream PLLs are: a PLL cell in an ASIC that receives an external reference signal, a PLL-based timing module that generates timing signals by multiplying an external reference, and a zero delay buffer used on a memory module to buffer the clock signal and provide the correct timing to latch the data. Although these devices may work properly with very stable reference inputs, if they cannot track a dynamically changing input signal then the output timing signals will not be synchronized to the system timing.

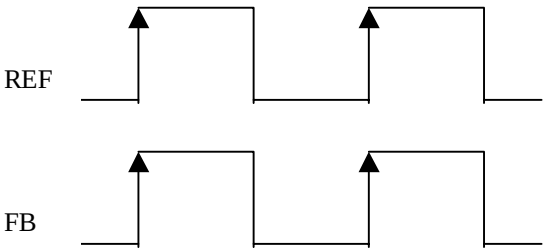

Figure 4. (a) Ideal Tracking Skew

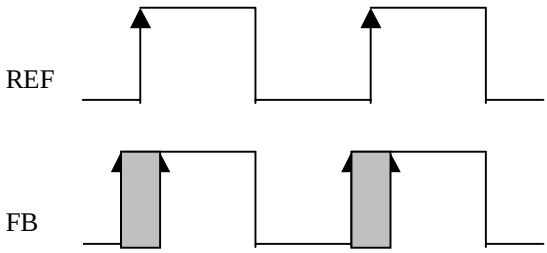

Figure 4. (b) Realistic Tracking Skew

Figure 4 shows a comparison of an ideal and realistic downstream PLL performance when tracking an SST reference signal. In *Figure 4(a)* the reference (REF) and feedback (FB) inputs to the phase detector are perfectly aligned and as the reference frequency changes the PLL loop responds instantly and the feedback signal perfectly "tracks" the reference in phase and frequency. In *Figure 4(b)* the real-world PLL has a non-zero response time and as the reference frequency changes the feedback signal will be slightly offset in phase as the loop responds. If we assume the reference frequency slightly increases then the feedback signal will slightly lag in phase until the loop responds. That is, the change of the reference input is detected by the phase detector, filtered, applied to the VCO, divided by the feedback divider and then applied to the feedback input of the PLL. The tracking skew is shown as both leading and lagging the reference input for two reasons: (1) because we are assuming the reference frequency is modulated around a stable center frequency, and (2) because at the scale of measurement we are concerned with here, pico-seconds, there is a statistical nature to at least a portion of the measurement. A typical measurement is made by triggering a very high-bandwidth digital oscilloscope on the reference signal and observing the feedback signal on a second channel. Using the infinite persistence mode of the oscilloscope, one is able to accumulate many traces of the feedback signal relative to the reference and measure the "accumulated tracking skew." Accumulated tracking skew is actually a combination of phase skew due to the PLL loop response and all sources of phase noise or jitter that are affecting the device, including input jitter.

When a spread spectrum timing signal is used as the reference input for a downstream PLL it must be able to track the frequency modulation without excessive tracking skew. Cypress uses the term Spread Aware™ to describe a PLL device that can properly track a spread spectrum timing signal. Spread Aware timing devices are able to follow the frequency modulated input signal with minimal tracking skew. Therefore, the system designer is able to maintain the best possible timing margins for latching data. In addition, the downstream subsystem will be compatible with systems that utilize SST technology. Timing devices that are not Spread Aware will either generate excessive tracking skew, thereby eroding the system timing budget, or not lock to the spread spectrum reference signal at all, resulting in system failure.

Figure 5 shows test results for two Zero Delay Buffers (ZDB) using spread spectrum timing signals as reference inputs. In both cases the oscilloscope is triggered on the reference input signal and the buffer output signal is captured using infinite persistence mode. The traces show the accumulated tracking skew which results from the PLL's loop response to the modulated input signal and the phase jitter of the buffer. *Figure 5(a)* is the output signal of an Cypress Spread Aware Zero Delay Buffer which is able to track the modulated input and has about ±80 ps of accumulated skew. *Figure 5(b)* is the output signal of a competitor's buffer which is obviously not able to track the spread spectrum modulated input.

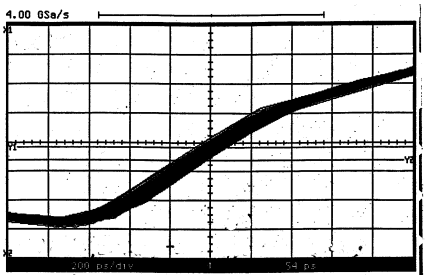

Figure 5. (a)IC WORKS Spread Aware Zero Delay Buffer

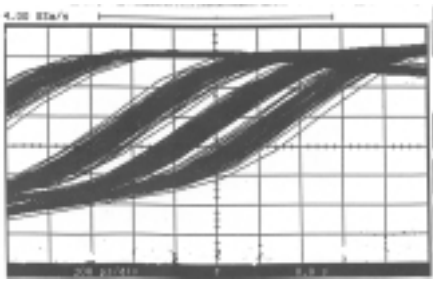

Figure 5. (b)Example of a ZDB That Cannot track an SST Input

Conclusion

The ability of a Spread Aware PLL to track a modulated reference input is based on the characteristics of the circuits shown in the PLL block diagrams. The characteristics which determine the overall loop response are: Phase Detector gain, VCO gain, feedback divider ratio and loop filter response. PLL designers strive for the best possible loop characteristics for a given application. They must balance many design choices, including loop bandwidth, response time and phase jitter characteristics. Therefore a PLL perfectly designed for one application may not work at all in a different system. Cypress can provide easy-to-use simulation tools so that PLL designers and users can simulate the response of a PLL to a spread spectrum input signal.

Cypress's family of Spread Aware Zero Delay Buffers are specifically designed to receive a spread spectrum modulated input signal. The PLL characteristics of Spread Aware devices will track the frequency modulation on the input signal with minimal accumulated tracking skew. Therefore, spread spectrum modulation present on the ZDB input signal will also be present on the output signals. This will reduce the EMI emissions of the system. In addition, since the PLL tracking skew has been minimized, the system designer will have the benefit of the greatest possible timing margins.

Sidebar – "PLL Basics"

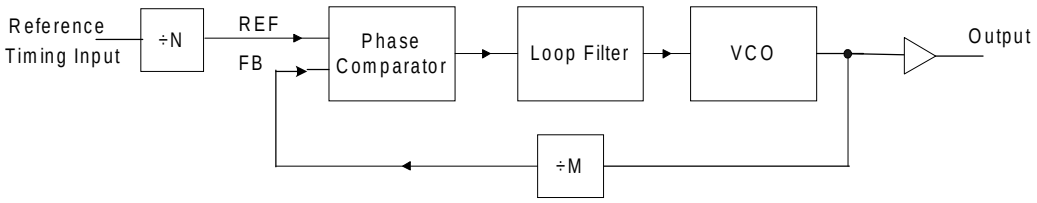

Simplified Phase Lock Loop (PLL) Block Diagram

Basic PLL Operation – non-varying reference frequency

One of the most common building blocks in timing circuits is the phase locked loop (PLL). As shown in the simplified block diagram above, a PLL consists of a Phase Detector, Loop Filter, Voltage Controlled Oscillator, and Feedback Divider. The function of the Phase Detector is to compare the phase and frequency of the reference (REF) and feedback (FB) input signals and produce a correction voltage that will be filtered and applied to the Voltage Controlled Oscillator (VCO). If the feedback signal is slower than the reference signal then the correction voltage will drive the VCO to a higher frequency until the reference and feedback signals are matched. In the alternate situation, when the feedback signal is faster than the reference signal, the correction voltage will reduce the frequency. The VCO frequency will settle at a value M times the Reference frequency because of the divider between the VCO and the Phase Detector. This is the basic technique that is used for frequency multiplication with PLLs. For example, if the phase detector reference is a 1-MHz signal and the feedback divider (M) is set to 16, then the VCO frequency will settle at 16 MHz. An extension of this technique is to add another divider between the input signal and the phase detector, shown as the optional divide-by-N in the simplified block diagram. The operation of the loop is the same as described above except that the phase detector reference signal is now the input signal divided by N. This will allow much greater variation in the frequency of the output signal because both N and M can be changed. The output frequency will now equal M/N times the input frequency.

Basic PLL Operation – modulated reference frequency

As described above, the basic operation of the PLL is to match the phase and frequency of the reference and feedback inputs to the Phase Detector. When the positive edges of these two signals are aligned the loop is in a "locked" condition. As long as the input frequency remains stable, the VCO frequency will remain stable. However, if the input frequency changes, then the VCO frequency will change also. In the example above, with the PLL configured as a 16x frequency multiplier, if the reference frequency is doubled from 1 MHz to 2 MHz then the VCO frequency will double from 16 MHz to 32 MHz. If the input frequency is dynamically changing then ideally the VCO frequency will follow or "track" the input and the phase detector inputs will remain aligned. When the loop is unable to track the input signal then the phase detector inputs become misaligned, resulting in a phase offset between the reference and feedback inputs to the phase detector. This phase offset is called "tracking skew." Tracking skew is the phase offset due to the inability of the loop to follow a changing input and is different from static phase offset that may be present even when the reference frequency is not varying.

Sidebar – Spread Spectrum Timing Glossary

Modulation Width or Spread Percentage: The relative variation in the instantaneous output frequency resulting from spread spectrum modulation (typically 0.5% to 4%).

The Spreading Signal: The low frequency signal which is injected to modulate the output.

Modulation Profile: The wave form of the spreading signal.

Modulation Frequency: The frequency of the spreading signal (typically 30 kHz to 50 kHz).

Centerspread: The nominal output frequency is specified, and the spreading results in instantaneous frequencies both above and below the nominal.

Downspread: The maximum output frequency is specified, and the spreading results in instantaneous frequencies at or below the maximum specified.

Tracking Skew: The phase offset of a PLL resulting from a time-varying reference input.

Accumulated Tracking Skew: The total measured phase offset due to tracking skew and phase jitter, including input jitter.

Downstream PLL: A PLL that receives a reference timing signal from another PLL-based device.