

Clock Recovery's Impact on Test and Measurement

Application Note

Introduction

Clock recovery plays a significant role in making accurate test measurements, whether incorporated into the test setup or as part of the device under test. As most gigabit communication systems are synchronous, the data within them are timed against a common clock. Whether traveling across inches of circuit board or traversing continents on optical fiber, the relationship between the data and the clock they were timed against can become disturbed. By extracting clock directly from the data, signal regeneration at the receiver can be achieved correctly.

It is important to note that receivers typically improve the incoming data signal before passing it on. A decision circuit in the receiver retimes the data and squares up the pulses. This process depends on a clock signal that is synchronous with the incoming data. Clock recovery within the receiver achieves this goal, provided the retiming clock moves in the same way and at the same time.

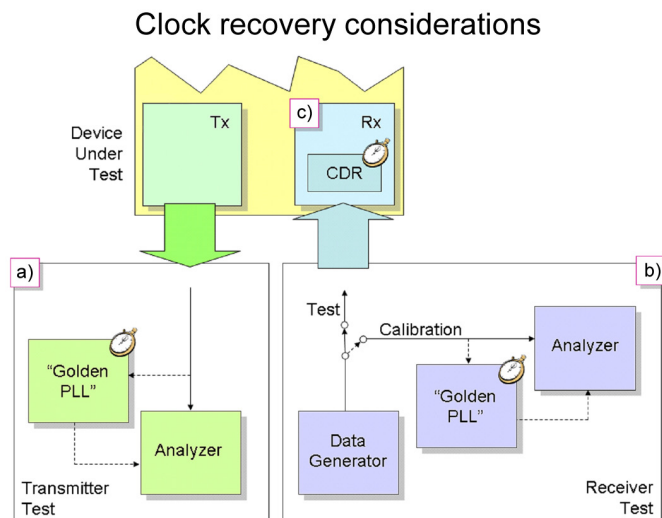


Figure 1. Clock recovery or 'Golden PLLs' can appear in several places in a test environment.

PLL-based Clock Recovery

Different architectures exist for achieving clock recovery. The most common methodology used in measurement equipment is based on a phase-locked loop (PLL). The recovery circuit is used to derive a clock that is synchronous with the incoming data, dependent upon seeing transitions in the data. For data segments with runs of identical bits, the PLL must remain locked. The loop gain has the most significant effect on the loop bandwidth; any filtering within the loop filter typically has a secondary effect. It should be noted that the transition density of the incoming data affects the amount of energy entering the loop and, therefore, the loop characteristics. The consequence is that loop bandwidths in compliance testing can change depending on the transition density of the pattern chosen.

The system transfer function performs a low-pass filtering operation on the phase modulation of the input signal, while the error-response transfer function performs a high-pass filtering function. The loop tracks input phase modulation within the loop bandwidth while failing to track phase modulation outside the bandwidth. This enables the loop to track low-frequency jitter and ignore high-frequency jitter outside the loop bandwidth of the PLL.

One measure of the jitter tracking characteristic of a PLL is the loop bandwidth (LBW), often measured as the point where the transfer of 'jitter in/jitter out' is -3 dB. However, this is not the only way that loops are defined.

A wide LBW improves jitter tolerance, while a narrow LBW removes more jitter from the recovered clock, which is beneficial to downstream synchronizers but impairs jitter tolerance. Although a wide LBW seems ideal, it is usually cost and technology dependent. Wide LBWs also contribute more noise or random jitter. Current LBWs used in measurements are typically in the range of 1 to 10 MHz.

The Ups and Downs of Clock Recovery

It is important to note how clock recovery is used in measurements and what could go wrong. On the transmitter testing side, for example, clock recovery is often required for two primary reasons — either because no clock signal is available to act as a test equipment trigger, or because standards require jitter measurements to be made with a specific LBW (see Figure 1, part a). The intention of the latter is that a system receiver, like a BERTScope BSA Series, will contain clock recovery that tracks out some of the incoming jitter, so transmitter test-ing should only be concerned with high-frequency jitter that a receiver would not track. See Figure 1.

Thus, for signals under test with jitter components near the LBW of the clock recovery, an incorrectly set LBW can result in inaccurate jitter measurements. Sometimes standards imply use of clock recovery in testing by either referring to a 'golden PLL' or by specifying measurement of jitter 'after application of a single-pole, high-pass, frequency weighting function that progressively attenuates jitter at 20 dB/decade below a frequency of $((\text{bit rate})/1,667)$ '.

Spread-spectrum clocking (SSC) spreads the energy of the clock (and the data) over 0.5% of the frequency band, lowering the average power at a given frequency in the spectrum. This helps products to comply with regulatory requirements for radiated and conducted emissions. For a receiver to successfully track out SSC, it must be able to track the modulation, including its harmonics, to avoid eye closure. If the loop response fails to sufficiently track SSC or there is an incorrect delay between the clock and data path, the test eye will be smeared closed.

Incorrect peaking (the region near the LBW where the jitter output of a clock recovery device can be greater than the jitter input) may amplify the amount of jitter measured. Additionally, trigger delay in the test equipment relative to the incoming data signal can cause incorrect amounts of jitter to be measured. For example, a fixed delay in the measurement system could cause additional apparent jitter to be measured. The additional jitter magnitude depends on the jitter frequency relative to the amount of delay.

At the receiver end, clock recovery can appear in either the device under test or as part of the test equipment calibration procedure. In the device under test, clock recovery is frequently present, and usually exercised in testing (see Figure 1, part b) using stress and sinusoidal jitter. In the case of sinusoidal jitter, testing is typically to a template whereby more jitter is applied at lower modulation frequencies and/or less jitter at higher frequencies.

Problems would include an incorrectly designed LBW in the receiver that could cause a jitter tolerance template to fail. Incorrect slope to the tracking response may not follow SSC accurately enough, resulting in a test eye that is smeared shut and causing bit errors to occur.

Clock recovery is frequently used in the test equipment setup and calibration of a receiver's jitter tolerance, or stressed eye, signal (see Figure 1, part c). Sinusoidal jitter is usually set to a frequency above the LBW of the clock recovery during calibration. However, an incorrect LBW may cause the wrong amount of stress to be set, resulting in either under- or over-stressing the device being tested and increasing the likelihood of customer rejection or yield issues, respectively.

From all these situations, it is easy to conclude that the LBW setting is critical and can have a significant impact on the observed jitter in measurement. Varying the loop bandwidth can give an indication of the jitter spectrum. Testing with very narrow LBWs can show all the jitter a transmitter under test is creating. Testing with very wide LBWs, on the other hand, indicates only the jitter that a transmitter produces that its intended system receiver is not able to filter out with its own PLL. Typically, the latter clock recovery is specified in compliance testing. A system designer is mainly interested in jitter that is beyond the capabilities of the receiver to track.

Distributed Clock Schemes

Not all systems derive their timing from the data stream. Some, such as PCI Express and fully buffered dual in-line memory module (DIMM), use a distributed clock routed to each end of the communications link for timing data. PLLs are used on the transmit and receive ends to multiply up the reference clock.

Typically, the distributed reference clock will have a certain amount of jitter; for instance, from the phase noise of the originating crystal. It could also have SSC. The clock is multiplied up within each IC and used to clock transmit and receive functions. Each PLL will have a loop response, and if their behavior is identical, jitter on one should be tracked exactly by the other, such that the receiver sees no net effect. However, reality tends to be more complex.

Even for devices manufactured with the same design, fabrication process, and manufacturing lot, it is virtually impossible to get identical loop responses. As it is also difficult to ensure identical path lengths between and within ICs, the equivalent of trigger delay is also apparent in the receiver jitter, causing it to see more jitter.

Embedded Clock Schemes

Embedding the clock into data is a common method of ensuring that accurate recovery of the transmitted data stream is achieved at the receiver. Once achieved, however, there is an issue of having a system running at one clock rate while an incoming bit stream is running at a slightly different rate. Somehow, the data must be reclocked to match the receiver system.

In some architectures, particularly SONET/SDH, significant effort is made to keep all the clocks in the system as closely matched as possible. This is achieved by distributing a highly accurate system clock based on a global positioning system (GPS) or local clocks based on rubidium or similar standards.

Other architectures assume more dissimilarity in clock rates to keep costs and complexity down. In any case, eventually the system must deal with any mismatch. This is typically done by waiting until the difference becomes more than one bit, or one frame, and then inserting or deleting bits or characters. Often the system protocol will insert characters, known as fill-words, that can be sacrificed at the receiver. At other times, the protocol will allow the receiver to insert characters of its own, if required, without disturbing the meaning of the data.

The addition or deletion of these characters can significantly affect testing. Protocol-based test equipment is usually set up to deal with inserted or deleted characters and can still recognize the underlying information. However, physical layer test equipment is sometimes more limited, requiring patterns to conform exactly to an unchanging, known sequence that repeats. Extra or missing bits cause the equipment to conclude that errors have occurred.

Data pattern changes can also occur in a system's management of baseline wander, the way a system can be thrown off by AC coupling and long runs of identical bits, causing the average signal voltage to drift until bit errors occur. In this case, protocol schemes often have two versions of each valid character and decide to send the one that most effectively counteracts any baseline wander or running disparity. The protocol intelligence at the receiver has no problem recognizing either version as being correct, but again, this violates the need for an unchanging bit pattern in some test equipment.

Some test equipment can make parametric measurements without the need for repeating patterns. This can be very effective at examining physical layer problems, but will be blind to protocol mistakes. Receiver bit errors that are cleaned up and retransmitted as healthy bits may also be missed, despite those bits being flawed.

Using loop-back testing, a signal sent into a receiver is looped back to become the output of the transmitter. Still, it does not always follow that the data will be identical, because mismatches in clock rate can cause fill-word changes that might upset test equipment. A solution in some circumstances is to create a test situation where the transmitter and receiver clock domains are absolutely identical, negating the need for domain rate matching. Schemes exist that use instrument clock recovery to create a clock signal at the exact rate of the test device output, using this signal to generate a test signal for loop-back testing.

As clock recovery becomes increasingly common in more systems and test setups, its effects on measurements must be considered. Many outside influences can disturb the relationship between data and how it is clocked. By understanding the relationship between both, more useful and accurate measurements are achievable.

Contact Tektronix:

ASEAN / Australasia (65) 6356 3900
Austria* 00800 2255 4835
Balkans, Israel, South Africa and other ISE Countries +41 52 675 3777
Belgium* 00800 2255 4835
Brazil +55 (11) 3759 7600
Canada 1 (800) 833-9200
Central East Europe, Ukraine and the Baltics +41 52 675 3777
Central Europe & Greece +41 52 675 3777
Denmark +45 80 88 1401
Finland +41 52 675 3777
France* 00800 2255 4835
Germany* 00800 2255 4835
Hong Kong 400-820-5835
India 000-800-650-1835
Italy* 00800 2255 4835
Japan 81 (3) 6714-3010
Luxembourg +41 52 675 3777
Mexico, Central/South America & Caribbean 52 (55) 56 04 50 90
Middle East, Asia and North Africa +41 52 675 3777
The Netherlands* 00800 2255 4835
Norway 800 16098
People's Republic of China 400-820-5835
Poland +41 52 675 3777
Portugal 80 08 12370
Republic of Korea 001-800-8255-2835
Russia & CIS +7 (495) 7484900
South Africa +27 11 206 8360
Spain* 00800 2255 4835
Sweden* 00800 2255 4835
Switzerland* 00800 2255 4835
Taiwan 886 (2) 2722-9622
United Kingdom & Ireland* 00800 2255 4835
USA 1 (800) 833-9200

* If the European phone number above is not accessible, please call +41 52 675 3777

Contact List Updated 25 May 2010

For Further Information

Tektronix maintains a comprehensive, constantly expanding collection of application notes, technical briefs and other resources to help engineers working on the cutting edge of technology. Please visit www.tektronix.com



Copyright © 2010, Tektronix. All rights reserved. Tektronix products are covered by U.S. and foreign patents, issued and pending. Information in this publication supersedes that in all previously published material. Specification and price change privileges reserved. TEKTRONIX and TEK are registered trademarks of Tektronix, Inc. All other trade names referenced are the service marks, trademarks or registered trademarks of their respective companies.

09/10 EA/WWW

65W-26074-0

Tektronix[®]

