Please note that effective July 26, 2021, the former Infrastructure and Automotive business of Silicon Labs is now part of Skyworks' Mixed Signal Solutions. All Silicon Labs registered trademarks that may be contained herein remain the sole property of Silicon Laboratories, Inc. and are only for nominative descriptive purposes and do not represent any sponsorship or endorsement of such product(s) by Silicon Labs.

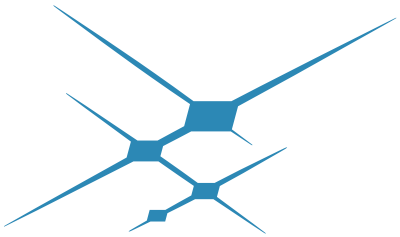

SKYWORKS®

skyworksinc.com

DSP-driven High-performance Clock Sources Radically Alter System Timing Architectures

Jerrell P. Hein and Mike Petrowski
Silicon Laboratories Inc., 7000 W William Cannon Drive, Bldg. 1, Suite 200, Austin, Texas 78735
jerrell.hein@silabs.com, mike.petrowski@silabs.com

1. Multi-protocol support drives low jitter clock diversity

As both business and residential customers demand greater access to high-speed network services, equipment suppliers are searching for ways to provide lost cost, multi-protocol interfaces that can be quickly provisioned for use. A prime example of this trend is the convergence of various data rates in the vicinity of 10 Gbps that can all share much of the same hardware and optics, even though each is governed by a different protocol as listed below in Table 1.

Table 1. Data rate as a function of protocol for 10 Gbps applications

Protocol	Data Rate (Gbps)	Forward Error Correction (FEC) Ratio	Line Rate (Gbps)	Typical Reference Clock (MHz)
SONET OC-192, SDH STM-64	9.95	-	9.95	622.08
G.975 (4 x OC-48 + FEC)	9.95	255/238	10.66	666.51
OTN OTU2 (G.709)	9.95	255/237	10.71	669.33
10 Gb Ethernet LAN	10.31	-	10.31	644.53
	10.31	255/237	11.10	693.48
10G Fibre Channel	10.52	-	10.52	657.42
	10.52	255/237	11.32	707.35

At the same time as performance and flexibility requirements of network interfaces are increasing, the expected levels of quality are also increasing, creating a need for more sophisticated and efficient test techniques to guarantee not only hardware functionality but design margin against the effects of aging, temperature variations, system noise and incoming signal degradation.

In contrast to rapidly changing system requirements to support multiple data rates and improved testability, the high-performance timing source technology used within these systems has failed to keep pace with the growing need for increased timing flexibility. Instead, timing subsystems remain tied to fixed frequency high-Q elements such as quartz crystals or surface acoustic wave (SAW) devices that meet the necessary performance requirements but lack the multi-frequency timing capability needed to support new services and test requirements. As a result, there is a growing need for new solutions to the problem of providing high-speed, low jitter clocks at the network interface.

2. Traditional multi-rate clock solutions

System designers have been able to provide a limited degree of data rate flexibility while still employing existing clock source technology. These designs are typically faced with the challenge of multiplying a low frequency network synchronization clock up to multiple high-frequency reference clocks as shown in Figure 1. Additionally, the timing subsystem must monitor the health of the network synchronization clocks while providing the capability to hitless switch between input references without causing phase transients on the transmit reference clock. To further complicate the timing subsystem design, sub picosecond (RMS) jitter clock requirements necessitate the use of jitter attenuating clock multiplier phase-locked loops (PLLs) of the type commonly constructed using discrete voltage-controlled SAW oscillators (VCSO) devices, phase detectors and loop filter elements. These PLL designs also need to support non-integer clock multiplication ratios needed for translation between base data rates and FEC line rates.

To support line cards with multiple data rates and FEC scaling ratios, designers rely on parallel instantiations of fixed multiplication ratio PLLs to produce the desired system clock rates. In these designs, the VCSO of each PLL is centered on the output clock frequency associated with each data rate and the RF multiplexer chooses only one of the PLL outputs at any given time. While this approach will support a small number of unique data rate choices, it does so at the cost of significant board area and component expense. Recently, some VCSO suppliers have introduced single packages that include two unique frequency VCSOs in order to reduce the board space penalty; however, this only provides incremental board space improvement and the solution still includes the cost of a separate resonator element for each unique output frequency/clock multiplication ratio combination.

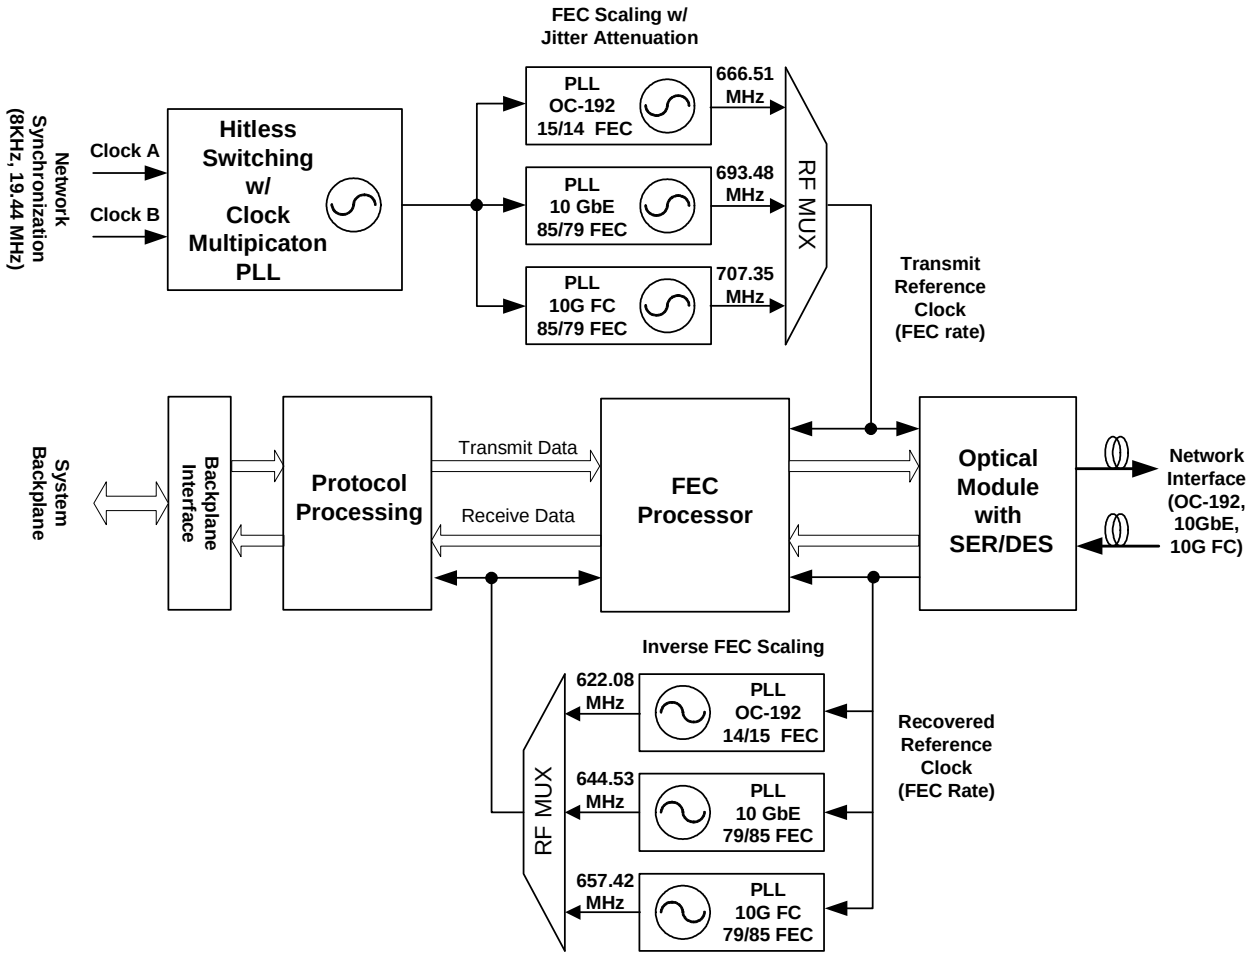

Fig. 1. Multi-rate 10 Gbps linecard using parallel fixed frequency PLLs

Alternatively, the system designer can specify a different bill of materials (BOM) for each data rate, creating a unique board type for each case. While this approach does not incur the board area penalty, as does the above parallel architecture, it does not allow for software provisioning of the service in the field and it increases the number of board types that must be kept in inventory. Also, neither approach provides any enhancement to the board test capability. System level timing architectures such as these have not changed much over time because of the lack of viable alternative technologies. The architectures have the appearance of being stopgap measures until more efficient solutions are found.

3. Digitally-controlled oscillator (DCO)

Recently, integrated circuit (IC) designers have started taking advantage of high-density, high-speed CMOS technology to develop digital signal processing (DSP) intensive clock source solutions that are both high performance and frequency agile. These DSP-based architectures (see Figure 2) use a low frequency resonator element (typically a quartz crystal) and a high-frequency on-chip VCO to produce a frequency agile high-speed, low-jitter output clock whose output rate is digitally-controlled and whose jitter performance equals that of traditional high-performance VCSOs. The resolution of the digital frequency control can be very fine, much less than one ppm, with a continuous tuning range of more than one GHz. Compared to the high frequency (>100 MHz), high absolute accuracy (<±20 ppm) and pulled (±20-100 ppm) resonators required in traditional high-performance VCSOs, these resonators can be very small and inexpensive because the reference resonator is low frequency (<40 MHz), has loose absolute frequency accuracy requirements (<±10,000 ppm) and is not pulled with changes in DCO output frequency. These resonators can be very small and inexpensive.

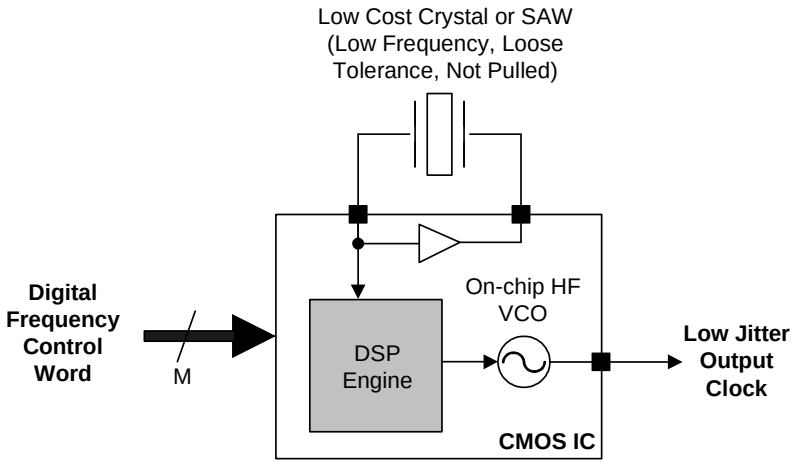

Fig. 2. Frequency agile high performance oscillator

4. New DSP enhanced PLL architecture

Utilizing the digital control interface provided by the DCO, a fully integrated digital PLL becomes possible that takes advantage of digital signal processing (DSP) algorithms as shown in Figure 3. In this DSP-based PLL architecture, the phase detector output is converted to digital format by a high-speed analog to digital converter (ADC). Following the ADC, all signal processing is done in the digital domain using high-speed DSP algorithms. The wide tuning range of the DCO (~15%) when combined with high-performance output dividers enables one PLL design to support a wide range of clock multiplication factors that would normally require multiple VCSO-based PLLs. In addition, the phase-noise performance of the silicon-based DCO is equivalent to that of fixed frequency VCSO alternatives, enabling narrowband loop operation for applications requiring jitter attenuation. The relative phase-noise performance of VCSO-based clock multiplier hybrids and a fully integrated DCO-based clock multiplier is shown in Figure 4. The biggest difference can be seen at high frequencies where the lower thermal noise of CMOS PLL yields improved jitter performance compared to the hybrid approaches.

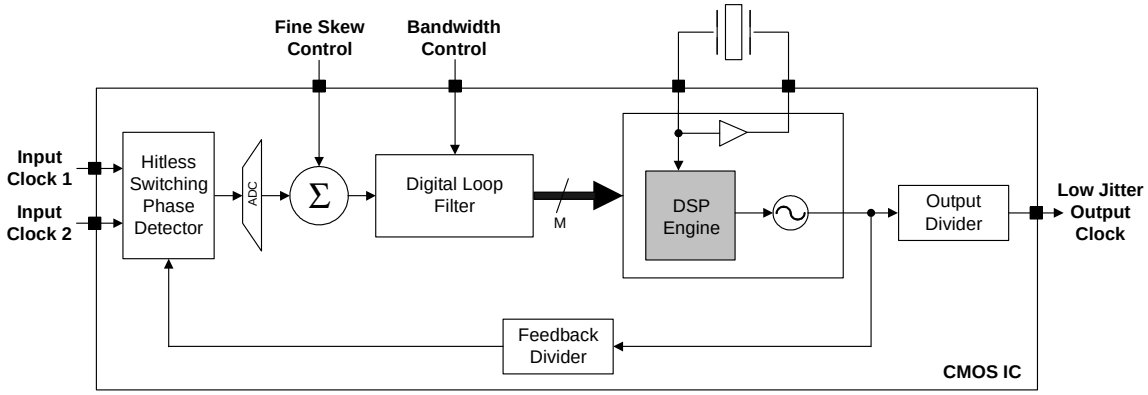

Fig. 3. High-performance CMOS PLL with embedded DCO

Integration of the PLL into a sub-micron CMOS process technology enables new digitally-controlled features such as hitless switching, clock multiplication, skew control and jitter attenuation. Digitally-controlled hitless switching virtually eliminates the output clock phase transient normally associated with switching the PLL input clock between two asynchronous clock sources, a function required in many optical port card architectures. This approach eliminates a cascade of narrow bandwidth PLLs that have traditionally been used to control phase transients during a clock switch. Programmable clock multiplication supports integer as well as non-integer scaling ratios, simplifying the translation to FEC frequencies or the translation between datacom and telecom data rates. Fine skew control allows adjustment of the input-to-output clock phase in picosecond increments, and digital bandwidth control allows static or dynamic modification of the PLL loop bandwidth in order to enhance PLL locking behavior and increase filtering of incoming clock jitter. These new DSP-based PLLs enable radical modifications in the system timing architecture for line cards providing multi-rate, multi-protocol support and also enabling improved test capability.

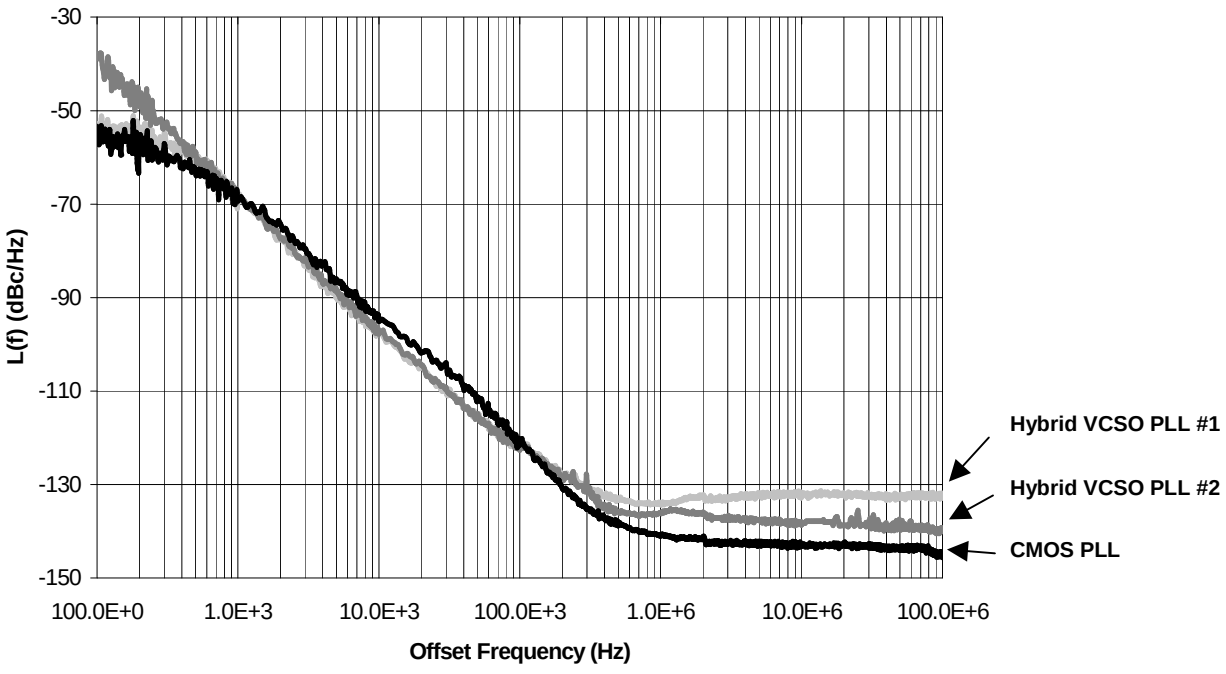

Fig. 4. 622 MHz Phase noise comparison of hybrid VCSO-based PLLs and CMOS-based PLL

5. System level simplification

The emergence of frequency agile, high-performance clock sources simplifies the timing subsystem in multi-rate, multi-protocol network interfaces. As an example, the system level timing architecture for the multi-rate and multi-protocol interface card described above (see Figure 1) can be dramatically simplified to the architecture shown in Figure 5. In this example, parallel banks of VCSO-based PLLs, RF multiplexers and hitless switching PLL circuitry can be eliminated and replaced with a single CMOS IC that is rate programmable via software control.

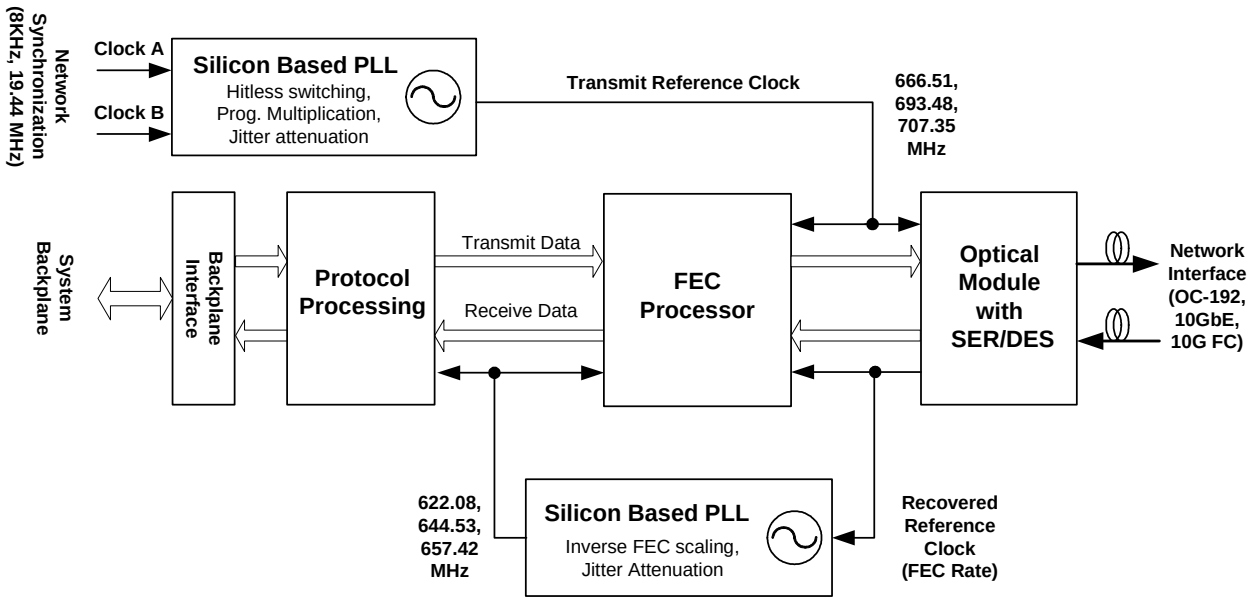

Fig. 5. Multi-rate linecard using frequency agile silicon-based PLLs

6. Manufacturing improvements

Independent of the need for multi-rate, multi-protocol capability, frequency and phase agile clock source technology greatly expands and improves the quality of board level testing. Figure 6 shows a typical board level test technique employed in the fixed-rate PLL system previously shown in Figure 1. Because the hardware can only operate at one frequency using the on-board timing components, the on-board clock sources are disabled during board test and external frequency programmable clock sources are used to drive the clock nets. Margin testing is performed by either raising or lowering the clock rate in small increments until hardware failure or by testing that the hardware operates properly at a predetermined amount above and below the nominal data rate. There are numerous disadvantages to driving the board clocks externally. For high rate signals, it is difficult to obtain adequate impedance matching between the external probe and the board trace, resulting in signal integrity degradations such as ringing or variable delay and amplitude. Also, probe points must be provided on the board that may compromise transmission line impedances and add undesirable capacitance to sensitive traces. Finally, additional test hardware is required to produce the test clocks.

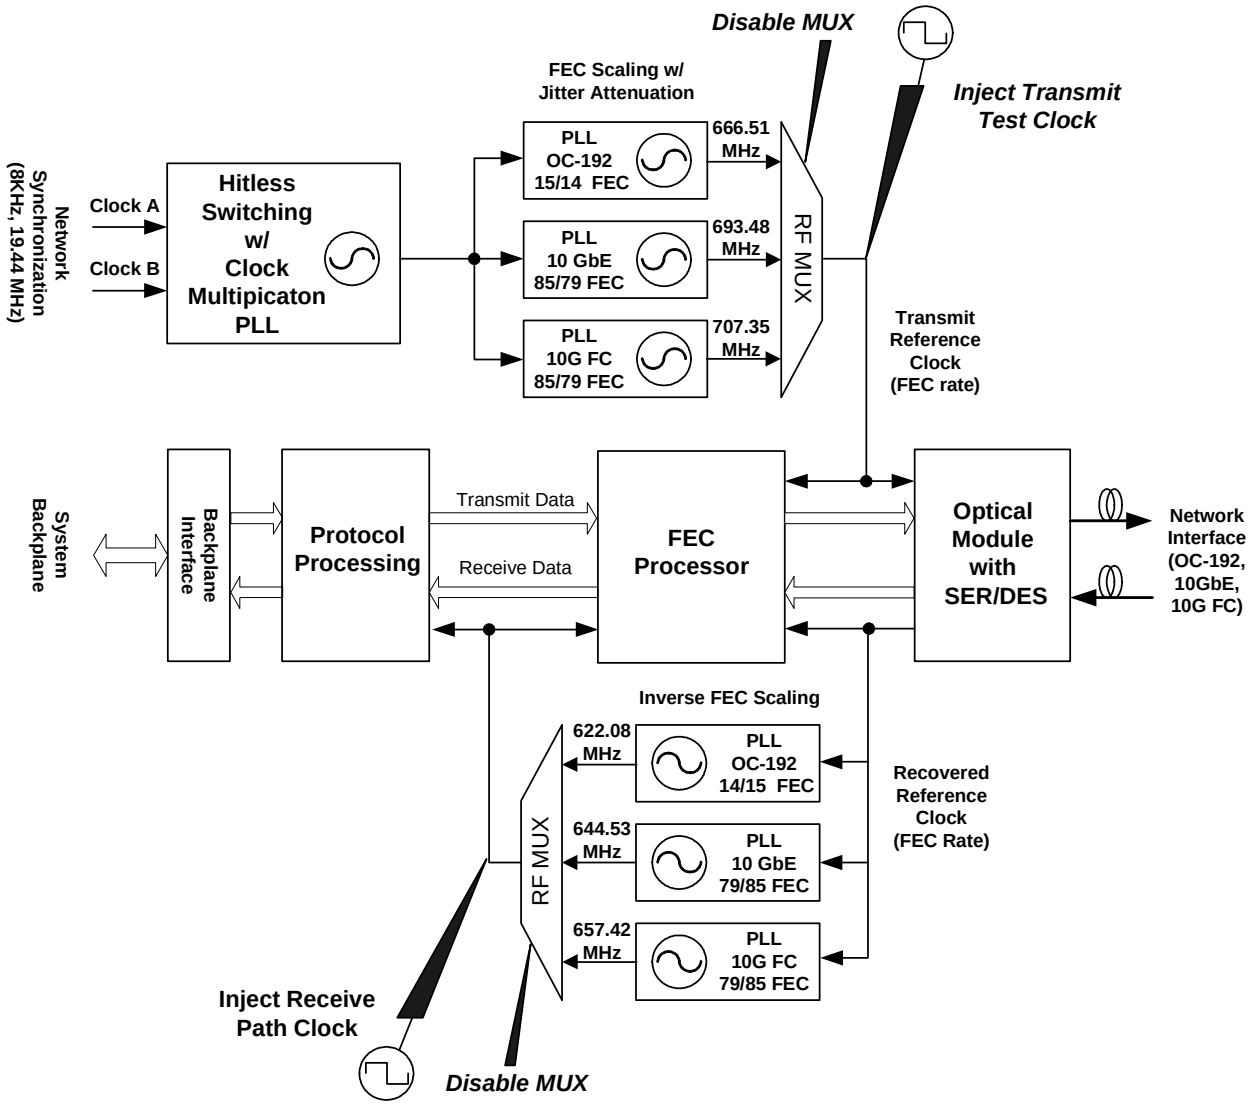

Fig. 6. Clock signal insertion for margin testing

With frequency agile clock sources on-board, external clocks do not need to be injected in place of the on-board clock sources, eliminating all of the above mentioned disadvantages. Signal integrity is maintained between normal operation and test mode, signal traces can be optimized for minimum capacitance and optimal transmission line continuity and no additional test hardware is required. Under software control, the output clock rates of the timing source ICs can be incremented or decremented in fine steps to determine the exact amount of frequency margin for the board under test.

Clock skew management is also supported using DSP-based PLL technology. Traditionally, system board designers estimate signal propagation times and size trace lengths to nominally maximize timing margins. However, once the board is manufactured, it is difficult to verify whether actual timing has been optimized, risking an additional circuit board revision and associated manufacturing delays. In addition, board dimension tolerances create variations in timing margin from board to board that cannot be compensated. While it is difficult to manage clock skew given variation in the manufacturing process, the software controllable fine latency adjustment available in DSP-based PLLs (see Figure 3) allows adjustment of delay between input and output clocks on a board by board basis. As a result, clock delay can be increased or decreased to determine the timing margin on each board, and the optimal delay for maximum operating margin can be programmed. This approach removes the risk of poor yield as a result of manufacturing variances or design miscalculation.

7. Looking forward

The conflict between the irresistible force of increasing system rate and protocol flexibility and the immovable object of fixed frequency timing sources is being resolved by the development DSP-based high-performance clock sources that are frequency and phase agile. Along with the system level benefits of software programmable data rates and protocols, board level testing will be improved, increasing manufacturing yield and timing margin. The net effect of this new technology will be lower cost and more easily provisioned high-speed services to network customers.