DesignCon 2007

Serial Protocol Compliance of an FPGA-Integrated Mixed-Signal Transceiver

Divya Vijayaraghavan, Altera Corporation

Ramanand Venkata, Arch Zaliznyak, Michael Zheng,
Steven Shen, Binh Ton, Lana Chan, Steve Park,
Chong Lee, Rakesh Patel, Richard Cliff,

Altera Corporation

CP-01022-1.0
January 2007

Abstract

The system-level protocol verification of a high-end FPGA with an embedded high-speed serial interface (HSSI) poses challenges that are comparable to and arguably exceed those encountered in ASIC-like verification flows. A single high-end FPGA device with embedded transceivers is designed to provide dedicated hard intellectual property (IP) support for a wide range of industry protocols and applications requiring high-speed serial I/O. This necessitates an efficient verification strategy that deviates from the traditional ASIC flow, leveraging the common aspects of serial protocols while addressing their nuances and maximizing reuse of verification IP in the various stages of the validation flow to achieve time to market with fully functional silicon. This paper describes a strategy devised and successfully utilized for the pre- and post-silicon protocol verification of an embedded 622-Mbps–6.375-Gbps serial interface in an FPGA. The strategy's evolution into the validation flow for next-generation protocols such as PCI Express Gen 2 is also discussed.

Author Biography

Divya Vijayaraghavan is a Senior Member of Technical Staff at Altera Corporation, specializing in developing serial I/O interfaces. She has 11 years of experience in chip design and verification. Divya received her MS degree in electrical engineering from the University of Texas, Austin and her BTech degree in electrical engineering from the Indian Institute of Technology, Madras. She represents Altera on industry standards committees.

Introduction

The embedded 622-Mbps–6.375-Gbps HSSI referred to in this paper utilized a common physical coding sub-layer (PCS) and physical media attachment (PMA) in hard IP to address the wide range of serial protocols and applications shown in Figure 1.

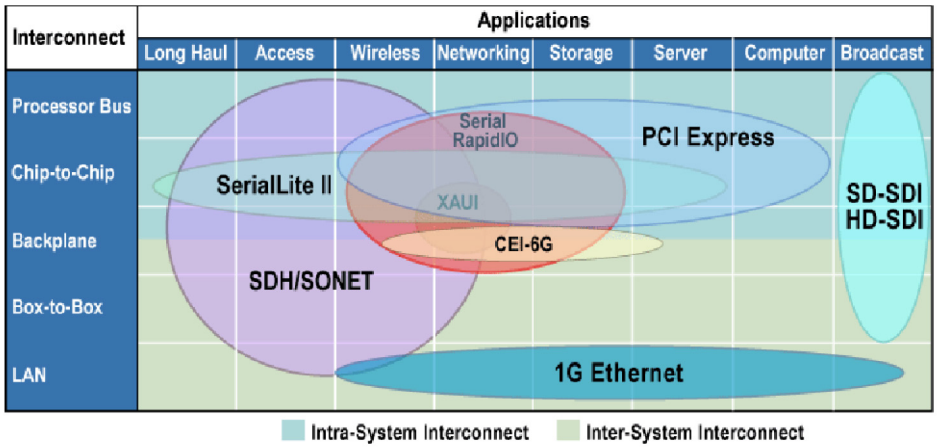

Figure 1. HSSI Serial Protocol Support

The high degree of configurability inherent in the hard IP, coupled with a modular architecture and the capability to enable or disable functionality as required, provided the customization required for each protocol and end application. Figures 2, 3, and 4 show simplified depictions of the HSSI configured in PCI Express/PIPE, Gigabit Ethernet, and XAUI modes.

A traditional verification methodology would utilize various distinct dedicated verification environments aimed at achieving the degree of functional coverage mandated by each serial protocol and configuration mode in the HSSI, which would naturally be detrimental to the time-to-market requirements of the device. An alternate, more efficient, verification strategy was therefore formulated that emphasized resource sharing and reuse throughout the verification cycle, while retaining the flexibility to cater to defining characteristics of each protocol.

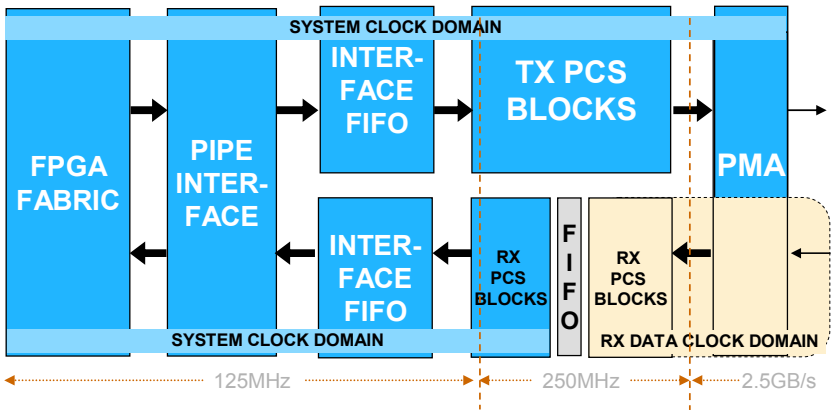

Figure 2. HSSI PCI Express/PIPE Configuration

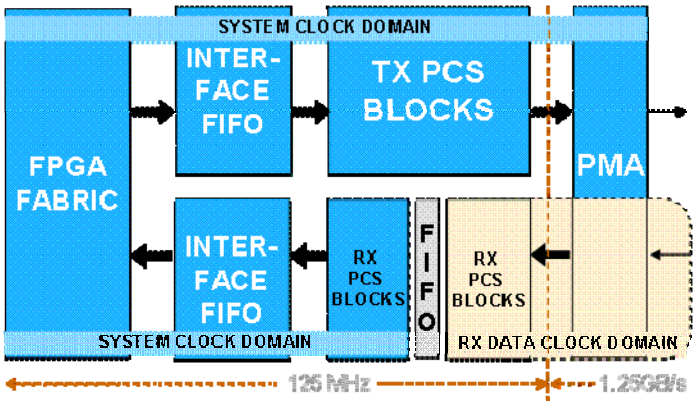

Figure 3. HSSI Gigabit Ethernet Configuration

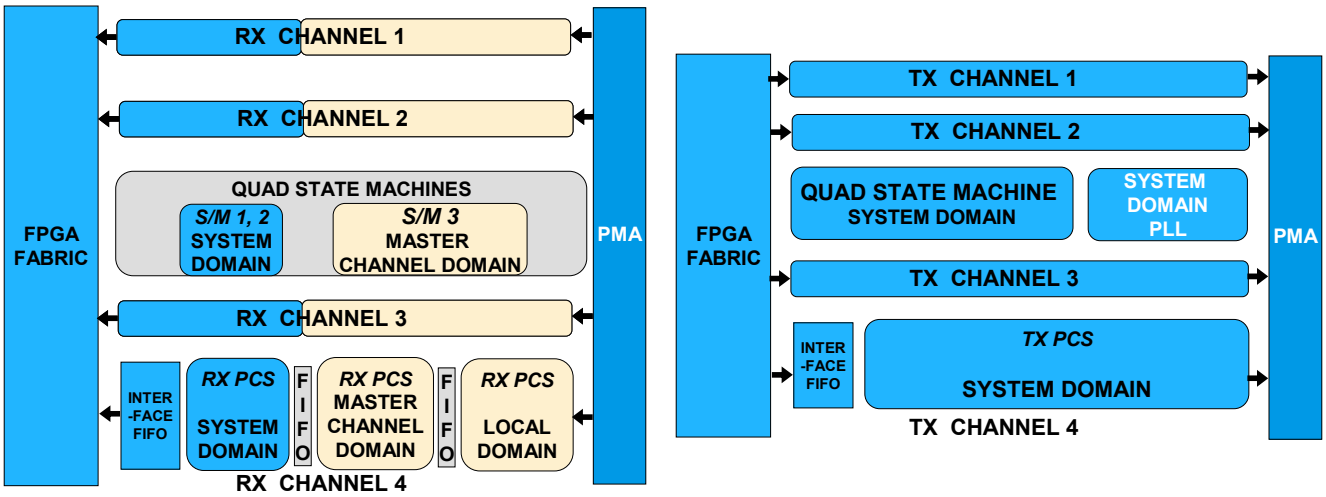

Figure 4. HSSI XAUI Configuration at 3.125 Gbps

Validation Criteria

The innumerable customization options provided by the HSSI were prioritized and classified into 33 modes targeted to current and future market requirements. These were further consolidated by protocol, and formed the basis for defining the validation requirements of the device.

A verification plan was created, comprising high-level, protocol-specific architectural requirements dictated by industry compliance checklists and lower-level requirements stemming from the implementation approaches in the design. Additionally, random testing requirements were identified, aimed at extending the coverage of architectural and implementation testing, targeting areas in the design that were envisioned to be complex and bug-sensitive. A combination of internal and third-party verification environments was chosen to cover the verification plan in its entirety. Selection criteria included the level of abstraction required (i.e., black box versus grey box testing), the ease of creation of specific verification scenarios, the ability to introduce randomness, the extent of reuse possible for silicon characterization, and extensibility to future generations of widely adopted serial protocols.

The third-party verification IP employed to achieve the protocol-specific architectural testing requirements of the HSSI provided off-the-shelf convenience, but brought with it the limitation of not being conducive to extensive resource sharing between protocols. This was illustrated by the dedicated PCI Express/PIPE and XAUI validation setups described in Section 3, which employed third-party bus functional models (BFMs) and MAC-layer IP from multiple vendors. In contrast, the in-house environments and test suites developed for implementation testing were sufficiently flexible to address the directed testing requirements of multiple protocols and configuration modes with minimal modification. The tests developed to verify the word–aligned state machine that provides the hysteresis associated with the byte synchronization function provided the coverage necessary for PCI Express and XAUI with minor user configuration changes, exemplifying this.

Both third-party and in-house verification environments and test suites were minimally adapted to form the basis to validate the functional requirements of next-generation protocols such as PCI Express Gen 2, though additional verification IP targeting the new functions was required to achieve complete functional coverage. An existing third-party PCI Express Gen 1 test suite was employed in conjunction with new PCI Express Gen 2 BFMs, and validated 60 percent of the functionality of PCI Express Gen 2. Eighty percent of the framework utilized to verify the clock compensation FIFO implementation for PCI Express Gen 1 and XAUI was reused to validate the FIFO buffer for PCI Express Gen 2. Enhanced multi-protocol test-benches with programmable parameters such as FIFO full and empty thresholds, latencies, read/write pointers, and insertion/deletion watermarks provided a high degree of flexibility to this in-house environment. Envisioned new functionality such as the PCI Express Gen 2 auto speed negotiation and electrical idle inference mechanisms could naturally not avail of significant test-bench reuse.

The verification plan and its random testing requirements in particular were repeatedly enhanced during the validation phase to address deficiencies uncovered by vector coverage metrics and the identification of bug-prone areas in the design.

Architectural and Implementation Testing

Independent black box testing environments were employed to ensure protocol compliance of the HSSI at the architectural level.

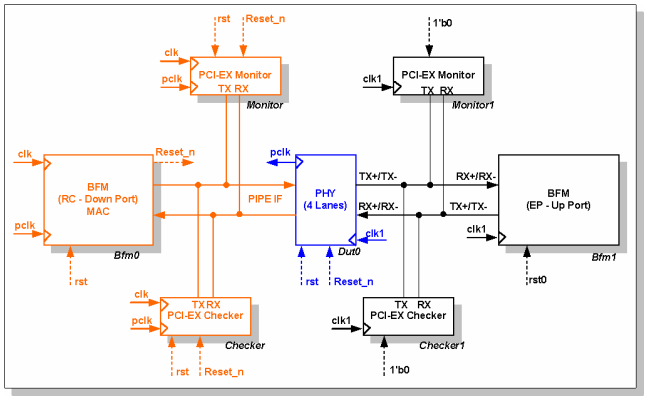

Figure 5. PCI Express Validation Platform With Single DUT

The validation environment in Figure 5 utilized a third-party PCI Express root complex MAC BFM interfacing to the device under test's (DUT) parallel PIPE interface and end-point MAC and PCS BFMs interfacing to the DUT's serial interface. The setup was reused with the BFM settings altered to validate the PCI Express/PIPE compliance of the DUT in an end-point role.

The third-party test suite accompanying the BFMs was enhanced to be event-driven, rendering changes in the DUT's interface timing and variations in the behavioral abstraction of the analog PMA transparent, as the design stabilized. Adequate coverage of corner cases, such as specific rate-match FIFO clock compensation scenarios, could not be achieved due to shortcomings in the configurability of the third-party PCS BFM, which prompted an additional setup employing back-to-back DUTs, as shown in Figure 6. This second validation platform enhanced the coverage of the rate matcher as required but introduced the risk of potentially masking non-compliant serial interface behavior.

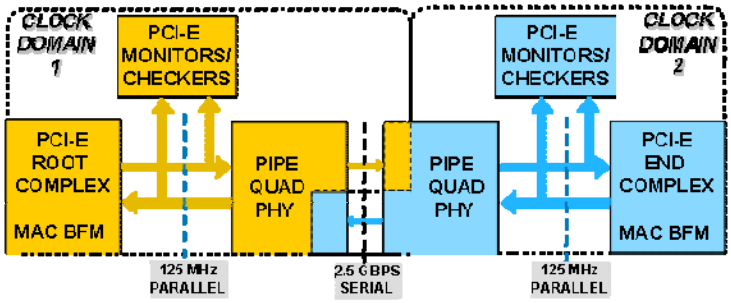

Figure 6. PCI Express Validation Platform With Dual DUTs

A PCI Express MAC IP Core from a second vendor was used as an additional PCI Express/PIPE verification source as illustrated by Figure 7, while an Ethernet MAC and its corresponding dedicated test suite were used to ensure XAUI compliance as shown in Figure 8.

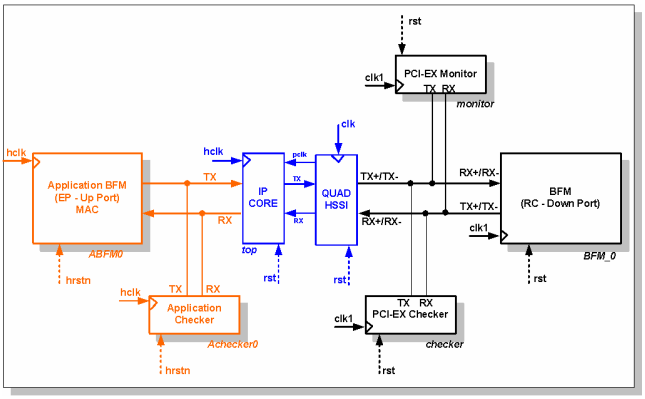

Figure 7. PCI Express Validation Platform With IP Core

Grey box testing of a more directed nature was employed to address the verification plan's implementation testing requirements, with common in-house test-benches addressing multiple protocols as discussed in Section 2 and illustrated by the module-level testing of the 8b/10b encoder and decoder. Directed testing was further utilized to create timing scenarios that were difficult to achieve in the environments discussed earlier, an example being the prioritization of error notification signaling in the PCS.

The optimal balance between black box and grey box testing was also determined by the extent of dedicated hardware support for each protocol within the HSSI.

Protocol Compliance of Analog-Digital Interface

The timing inaccuracies introduced by the Verilog behavioral abstraction of the PMA, utilized in Section 3, were addressed by the advance mixed-signal simulation (ADMS) environment in which the analog portion of the HSSI was almost entirely represented by its SPICE netlist. The 33 functional modes that formed the underlying basis of the verification plan discussed above were further consolidated into 12 ADMS testing templates, leveraging common settings in the configuration of the PMA.

The comparatively lengthy run times inherent in the ADMS approach were easily offset by the unique coverage of crucial aspects of the verification plan, such as clock and data transfer across the PCS-PMA interface, and dynamic reconfiguration that would enable the compatibility of the HSSI with next-generation serial protocols like PCI Express Gen 2. The ADMS setup also provided visibility into the timing associated with mandatory protocol-specific functions that spanned the PCS and PMA, specific examples being receiver detection and the N_FTS timeout requirement associated with low latency receiver L0s exit in PCI Express.

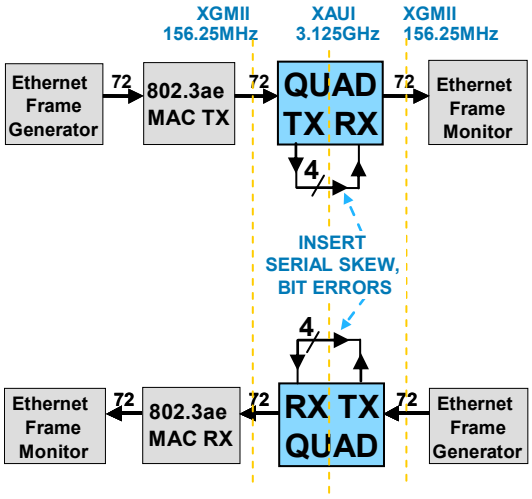

Figure 8. XAUI Validation Platform With IP Core

Emulation Platform Reuse for Silicon Characterization

The emulation platform for the pre-silicon protocol verification of the HSSI was configured based on the 12 ADMS templates, and designed to evolve seamlessly into the

characterization platform utilized for post-silicon protocol validation. The emulation flow comprised the synthesizing of the PCS portion of the HSSI on to an FPGA and the employing of external test equipment to configure the DUT and apply stimuli and checks to it, as evidenced by the flow diagram in Figure 9 and the sample laboratory setup in Figure 10.

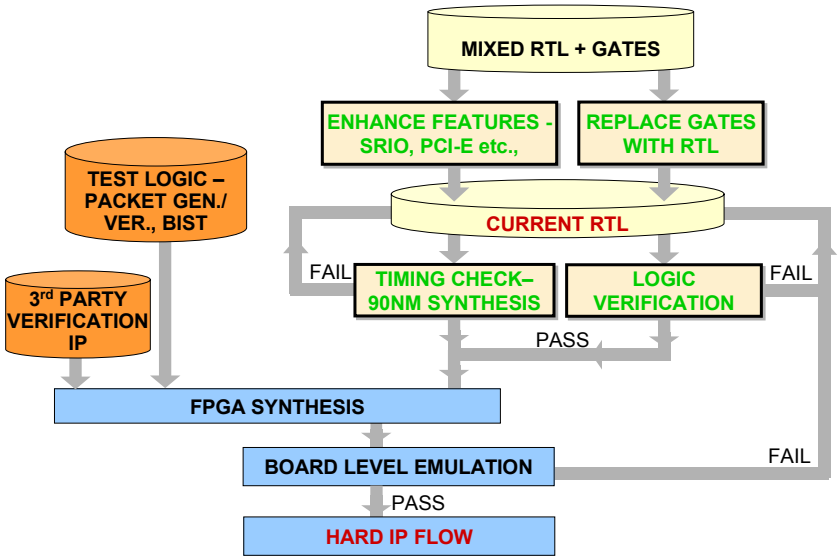

Figure 9. Emulation Flow Diagram

Emulation provided the advantage of a significant reduction in run time over simulation facilitating the verification of lengthy protocol-specific sequences of events such as the link width and lane number configuration scenarios and low-power entry and exit procedures driven by the PCI Express link training and status state machine. It also made feasible the pre-silicon injection of real-world aberrations such as clock-phase shifts and jitter, in order to hit the more directed clock-compensation corner cases that were common to multiple protocols and were a challenge to achieve in simulation. Emulation did however utilize a “virtual PMA” to abstract the high-level functions of the HSSI’s analog functionality and did not add value to the coverage of PCS-PMA interface timing, as it largely employed internal parallel loopback in its testing. Additionally, at-speed verification was prevented by the high resource utilization in the FPGA, which limited the throughput of the setup, as well as restricted the maximum link width of the DUT to four channels.

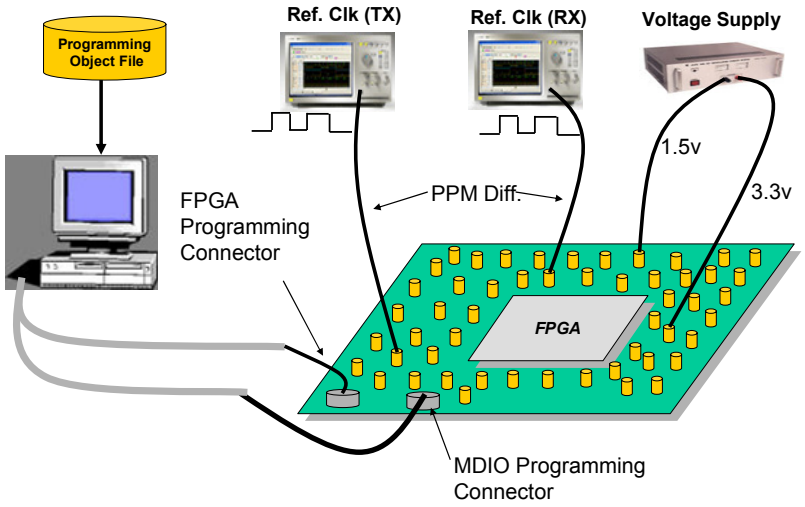

Figure 10. Emulation Laboratory Setup

Architectural and implementation-level protocol requirements were addressed by the transactor-based and ROM-based emulation platforms shown in Figures 11 and 12. The transactor-based platform employed an in-house packet generator and checker that could be configured to address the various framing rules of PCI Express, XAUI, and Gigabit Ethernet, providing directed function-based coverage to multiple protocols simultaneously. The more powerful ROM-based platform provided a higher degree of timing accuracy, utilizing stimuli and expected behavior captured from simulation and incorporating a state machine to enable self-checking.

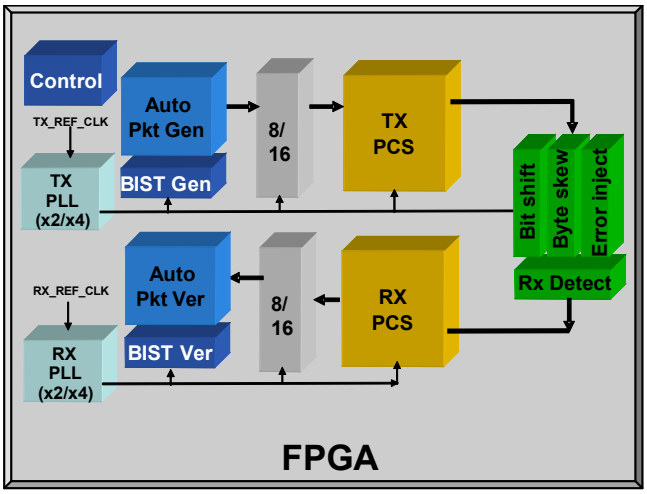

Figure 11. Transactor-Based Emulation Platform

Third-party serial protocol upper-layer IP was also utilized to provide system-level “real world” coverage of serial protocol functions such as the state transitions associated with recovery from an excessive bit error rate (BER) in PCI Express links and the XAUI static lane deskew function. The ROM-based and third-party setups did not inherently promote multi-protocol usage.

The smooth migration of the emulation setup to the at-speed characterization platforms employed for silicon bring-up is shown in Figure 13, which illustrates the replacement of the DUT on the FPGA, with actual silicon external to the FPGA in the transactor-based and ROM-based configurations.

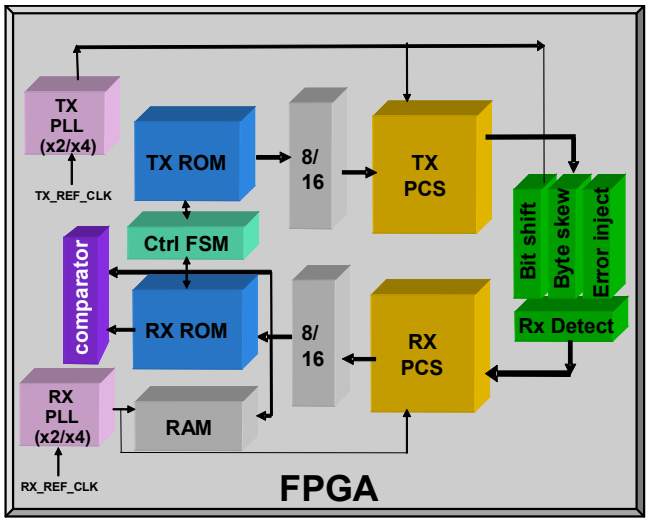

Figure 12. ROM-Based Emulation Platform

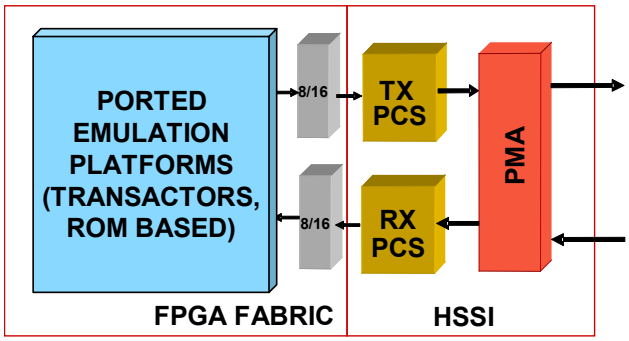

Figure 13. Silicon Characterization Platform

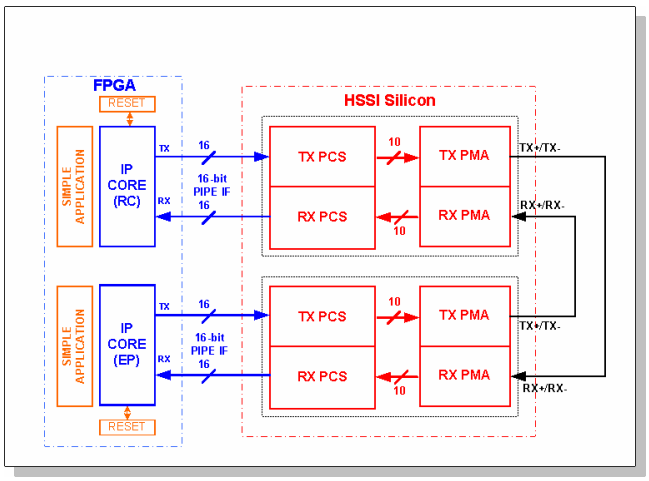

Figure 14. Third-Party IP-Based Characterization Platform for PCI Express/PIPE

Figure 14 shows a characterization setup employing third-party PCI Express upper-layer IP. The architectural and implementation coverage provided by simulation and emulation was further enhanced with the use of internal serial loopback incorporating the actual PMA and the ability to accommodate a link width of eight channels. This provided confirmation of the HSSI's analog-digital timing interactions and adherence to parameters such as transmit and receive inter-lane skew budgets and latencies in the various operational modes of the device. At-speed testing of the PCI Express Gen 2 PCS and PMA re-employed the Gen 1 characterization platform, utilizing 500-MHz Gen 2 and 250-MHz Gen 1 PCS frequencies and a largely common Gen 1-Gen 2 PCS design. Back-to-back characterization boards were used to create specific clocking PPM differences and reference clock settings to provide the final incremental directed coverage that had eluded the prior validation platforms.

Conclusion

The underlying theme of verification IP reuse and consolidation that formed the backbone of the efficient validation strategy described in this paper ensured the protocol compliance of the multiple functional modes of the embedded HSSI in minimal time. The aspects of the strategy that were inflexible and not conducive to resource sharing were discussed.

Further reduction in the duration of the verification cycle may be achieved through early bug detection with the use of assertion-based verification and formal verification methodologies to verify design intent prior to simulation. These and other optimizations will be vital to addressing the burgeoning serial protocol validation requirements of future generations of FPGA-embedded HSSIs.

Acknowledgments

The author would like to thank colleagues from the IC Design, Layout, CAD, Product Engineering, Software, Applications, and Product Planning teams at Altera whose valuable contributions are greatly appreciated.

References

[1] Malandain, D., Palmen, P., Taylor, M., Aharoni, M. and Arbetman, Y. "*An Effective and Flexible Approach to Functional Verification of Processor Families"*, Seventh IEEE International High-Level Design Validation and Test Workshop, 2002.

[2] Sauer, C., Gries, M., Gomez, J. I., Weber, S. and Keutzer, K. "*Developing a Flexible Interface for RapidIO, Hypertransport, and PCI-Express",* International Conference on Parallel Computing in Electrical Engineering, 2004.

[3] Kumar, J. and Pixley, C. "*Logic and Functional Verification in a Commercial Semiconductor Environment"*, International Conference on Application of Concurrency to System Design, 1998.

[4] Zaliznyak, A. et al "*Top-down Post-layout Mixed-Signal Design Verification Methodology for FPGA-embedded Multi-gigabit Transceiver",* DesignCon 2005.

101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
http://www.altera.com