

Architecture and Methodology of a SoPC with 3.25Gbps CDR based Serdes and 1Gbps Dynamic Phase Alignment

Ramanand Venkata, Wilson Wong, Tina Tran, Vinson Chan, Tim Hoang, Henry Lui, Binh Ton, Sergey Shumurayev, Chong Lee, Shoujun Wang, Huy Ngo, Malik Kabani, Victor Maruri, Tin Lai, Tam Nguyen, Arch Zaliznyak, Mei Luo, Toan Nguyen, Kazi Asaduzzaman, Simardeep Maangat, John Lam, Rakesh Patel
 Altera Corporation
 101 Innovation Drive
 San Jose, CA 95134

Abstract

The SoPC (System on a Programmable Chip) aspects of the Stratix GX™ FPGA with 3.125Gbps SERDES are described. The FPGA was fabricated on a 0.13um, 9-layer metal process. The 16 high-speed serial transceiver channels with Clock Data Recovery (CDR) provides 622-megabits per second (Mbps) to 3.125-Gbps full-duplex transceiver operation per channel. Another challenge described, is the implementation of 39 source-synchronous channels at 100Mbps to 1Gbps, utilizing Dynamic Phase Alignment (DPA). The implementation and integration of the FPGA logic array (with its own Hard IP) with the CDR and DPA channels involved grappling with SoC design issues and methodologies.

Introduction

Rapidly increasing data rates in a wide variety of applications ranging from telecom backplanes to HDTV video production environments is forcing a shift from parallel buses to serial interfaces. Among many advantages are cost savings resulting from reduced pin counts, an effective method of dealing with noisy system environments and elimination of clock to data set up / hold windows, which CDR facilitates (data is sent without any accompanying clocks.) FPGAs have come a long way from simple PLDs (1, 2) to the more complex high density Products of today (3, 4).

The High Speed Serial Interface (HSSI) Quad supports a range of standards and protocols – such as 10 Gigabit Ethernet XAUI, SONET/SDH, 1 Gigabit Ethernet (GIGE), PCI Express, SMPTE 292M, SFI-5, SPI-5, InfiniBand, Fiber Channel, and SerialRapidIO. Support for source-synchronous bus standards, include 10-Gigabit Ethernet XSBI, Parallel RapidIO, UTOPIA IV, Network Packet Streaming Interface (NPSI), HyperTransport™ technology, SPI-4 Phase 2 (POS-PHY Level 4), and SFI-4.

Fig. 1 is a SoPC application, which shows a generic architecture of a 3G base-station. This application utilizes several SoPC concepts. A soft embedded processor (for example, the NIOS® (4)) with customizable instruction sets can be used in the transceiver card. The transceiver cards receive and transmit data at 3.125Gbps using industry standard or proprietary protocols from and to the back plane via the channel cards.

The device architecture definition had to be scalable to reflect the varying bandwidth needs of SoPC (System On a Programmable Chip) applications. 4 to 20 CDR channels

at 3.125Gbps to 622Mbps allow for a combined bandwidth of up to 62.5Gbps. Also, up to 45 source-synchronous (clock accompanies data, but has arbitrary skew relationship with data) DPA channels at 1Gbps provide an additional 45Gbps bandwidth. The device family contains 10,570 to 41,250 logic elements and 330 to 544 I/Os.

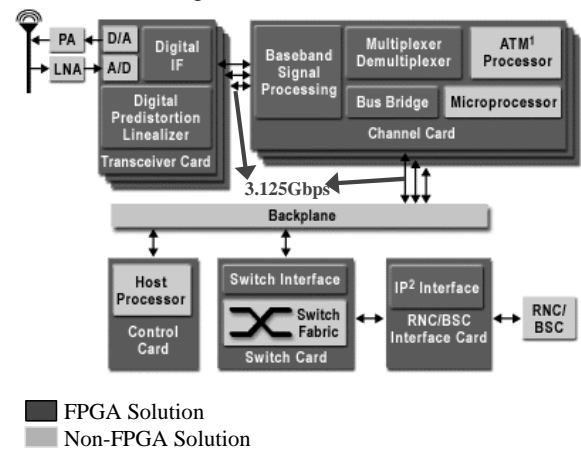


Fig. 1 3G Base-Station Architecture

(Portable blocks, whose architectures are predefined and layouts are “mostly fixed”, are referred to as Hard Intellectual Property or Hard IP blocks) Hard IPs integrated into the PLD fabric are: i) HSSI, ii) DPA, iii) High-speed DSP blocks that provide dedicated implementation of multipliers (at up to 250 MHz), multiply-accumulate functions and finite impulse response (FIR) filters, and iv) Up to 3.4Mbits of RAM (available without reducing logic resources.) This paper will discuss the 16 CDR / 39 DPA Channel device and will focus on the intergration aspects of these channels only.

Integration Challenges

The integration challenges can be broadly classified into 5 categories:

- A) Floorplanning - Chip level and Hard IP
- B) Architectural Integration
- C) Integration Methodology – Design & Simulation
- D) Layout Integration
- E) Package Integration

A. Floorplanning - Chip level and Hard IP

The first obvious hurdle was determining the floorplan of the Hard IP, at a time when both the FPGA and the Hard IP architectures were in flux. The only information available indicated that the packaging would employ a flip

chip array of bumps, with the I/O to bumps connections utilizing a redistribution layer. Another constraint was dictated by the Hard IP's high speed serial I/O, which required that the associated bumps should be located as close to the I/O as possible. A third constraint was imposed by a circumstance probably unique to Hard IP integration: the HSSI block's layout was required to be "dropped in" from a previously validated test chip. The Hard IP had to fit snugly into the area previously occupied by the I/O columns at either side of the FPGA, as shown in Fig. 2.

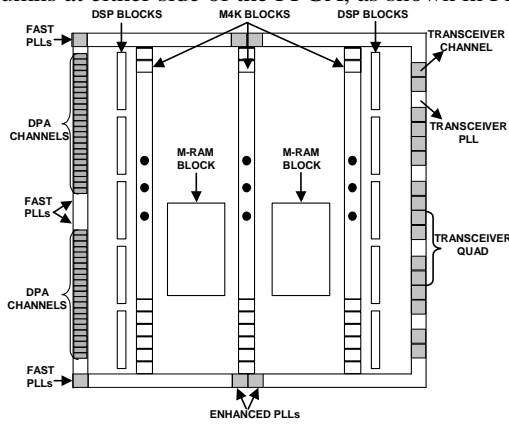


Fig. 2 Floorplan Overview

The next step was floorplanning to match the FPGA Logic Array resources such as LUT Elements (LE) and interconnect to the Hard IP. One distinguishing feature of the Hard IP blocks was that they were composed of several identical channels (a channel being made up of a receiver and a transmitter channel). Care was applied towards ensuring each channel's data and control traffic could efficiently be handled in the same LE row and had easy access to FPGA resources such as memory blocks. The result was a row to channel mapping, with 1 DPA channel (RX & TX) per LE row and 2 LE rows per HSSI channel.

B. Architectural Integration Overview

Even on it's own, the HSSI block faced issues confronting system designers. The SERDES usually is a separate discrete device and for obvious purposes. Analog designers with custom design techniques layout the PMA (Physical Medium Attachment), while the PHY (physical layer) or the PCS (Physical Coding Sublayer) is best left to the expertise of ASIC designers. But, this device integrated these disparate blocks as shown in Fig. 3.

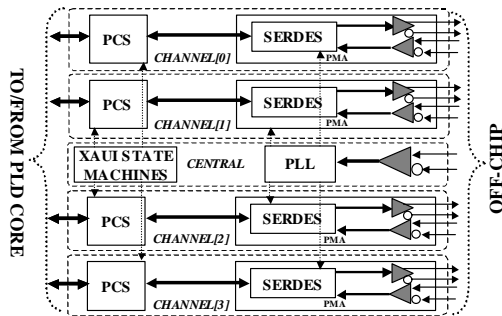


Fig. 3 HSSI Quad

The HSSI channels were organized into a Quad configuration. The XAUI standard (5) requires that 4 channels act in tandem to stripe 10Gbps data throughput across 4 channels, with XAUI (XGMII) state machines controlling the data path in the central block. This meant that one PLL had to control 4 channels, giving rise to the Quad configuration.

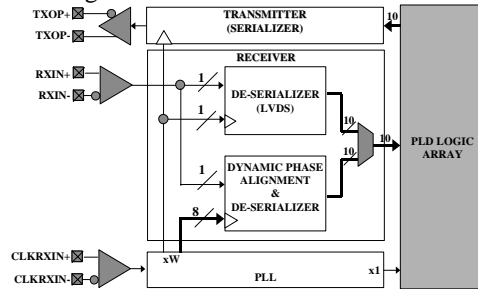


Fig. 4 DPA

In the DPA (Fig. 4), the PMA is composed of the receiver circuit and the transmitter circuit. The receiver circuit has both an LVDS scheme and the DPA. The PLL shown provides clocks for all the DPA channels. One of the 8 clocks is dynamically selected to provide the best data / clock skew – clock at the center of the data eye.

The interaction of the high speed HSSI and DPA (also referred to as Hard IP) with the slower speed FPGA Logic Array (also referred to as the FPGA Core) had to be carefully specified. Many other blocks in the FPGA were also designed for faster operation (with the embedded memory blocks operating at 350Mhz and the DSP blocks at 250MHZ). More architectural features to enable seamless integration are described in the sections below.

i) Clock Management

The HSSI digital block could be operating at up to 400MHz and thus, transfer of data and control signals to (from) Hard IP from (to) the slower speed FPGA core was a new challenge (Fig. 5).

An integral part of an efficient SoPC architecture is the flexibility provided with respect to clock selection and usage. Fig. 5 shows how the data transfer between the FPGA and Hard IP is done seamlessly by isolating the clock domains.

The clocks that are sent to the FPGA Core can be fed into a clock tree system (Fig. 6) that extensively covers the whole chip – DSP & memory blocks, I/O registers and LE




Fig. 5 Isolated clock domains

registers. Each Hard IP register at the interface is viewed by the clock trees as yet another FPGA Core register. This

significantly eases the FPGA fitting and routing software’s task managing up to 40 data bus (920 data signals) transfers. The device architecture provides up to 48 independent clock trees, with up to 8 on-chip PLLs.

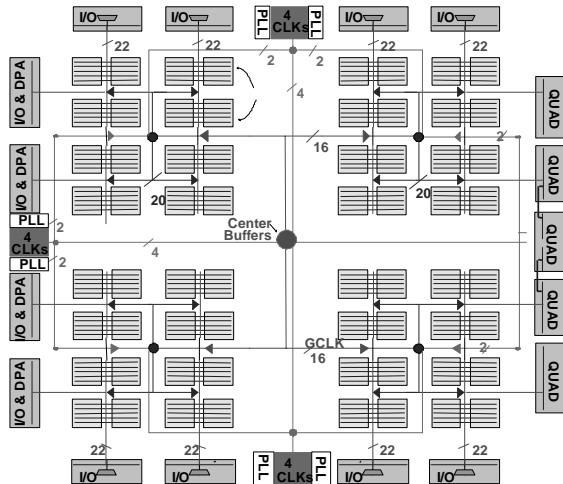


Fig. 6 PLLs and clock trees

Data buses can be optionally stepped up right after they enter the Hard IP and stepped down right before they exit the Hard IP as shown in Fig. 5. The key word is “optional”. Once again, SRAM configuration bits (CRAMs) were designed in for this purpose. These bits can also determine whether the recovered CDR clock or the PLL Clock should be divided before it is sent to the FPGA Core. The dividers and the muxes to bypass the demux are not shown in Fig. 5.

Reference clocks for the CDR PLLs and the transmit PLLs (in the Central area) could be chosen from a variety of sources: dedicated input pins and a special Inter-Quad clock network, whose inputs were from the PLD fabric.

ii) FPGA-Hard IP Control signals

The HSSI block and the FPGA Core exchanged several control signals, in addition to data. Since, the FPGA core could be operating at half the frequency of the Hard IP (Fig. 5), it was decided that control signals across the interface must be a minimum of 2 clock cycles long.

CRAMs from the FPGA were widely used to bypass any block or to choose different functionalities in the HSSI PCS. For example, SONET customers can bypass the 8B/10B Encode/Decode logic or GIGE customers can select the GIGE state machines and other special supporting blocks for GIGE. Pre-emphasis levels on the high speed output buffer could be dynamically changed during operation by control signals from the PLD fabric. To sum up, multi-standard support was emphasized through out the architecture.

C. Integration Methodology – Design & Simulation

An additional level of complexity was introduced, because the FPGA fabric was designed and verified in an entirely different design methodology from that of the Hard IP. This flow used a mixture of schematic entry and

in-house tools. The two Hard IP blocks followed separate design and simulation methodologies unique to each, because of the specialized nature of the blocks. Also, as is common with IP integration the individual IP design teams were geographically and (design-cycle wise) chronologically dispersed.

i) **HSSI:** The HSSI is a mixed signal block with ASIC and analog block components. The analog blocks were entered in a schematic entry tool. The digital blocks were specified in Verilog HDL and synthesized employing a complete ASIC methodology from synthesis to back-annotated crosstalk analysis.

The HSSI block was verified using a unique mixed signal simulation methodology. Purely analog blocks were modeled in Verilog. Digital elements (registers and combinatorial gates) in the analog schematics were replaced with equivalent Verilog models. This allowed realistic system level simulation taking into account the lock times of the PLLs and the CDR circuitry. In parallel, the integrity of the connections between the analog blocks were validated.

The second strategy used a commercial mixed signal simulation tool. This tool enabled Verilog test benches developed previously to be reused and enabled simulating a database made up of Verilog ASIC portions and Analog Spice netlists..

ii) **DPA:** The DPA was designed with standard cells, but still entered via schematics and timing verified with an ASIC-like verification flow. A commercial auto-router placed and routed the non-critical blocks, while the sensitive phase alignment blocks were custom placed and routed. The DPA schematics were converted into a Verilog netlist for functional verification purposes.

iii) Full Chip Cosimulation

The software team had to verify the accuracy of the bit map of the device - a one to one mapping of all CRAMs to its stated functionality. The team used an inhouse Design Tool - Quartus - to implement functions that exercised the thousands of CRAM bits. Quartus uses software models of all functional blocks in the chip. The IC Design team co-simulated (Fig. 7) Software’s bit mapping and vectors using the real “mixed Verilog / schematics” database. The outputs from the latter simulation had to match Quartus’s outputs.




Fig. 7 Co-simulation

The Verilog models of the Hard IP blocks along with the FPGA Core schematics were co-simulated in

Viewlogic Fusion, which has Verilog and schematic simulation engines exchanging data between each other.

D. Package Integration

The package supports a mixture of I/Os varying from low speed to 1Gbps (DPA) and up to 3.125Gbps (HSSI) signals. To support the high density of I/Os, flip-chip packages are employed ranging from 672 to 1020 balls. There are close to 200 high-speed traces that require advance SI evaluation. HSSI traces are extracted and analyzed using advanced tools such as HFSS and Ansoft to assure excellent package performance at operating frequency. Due to the large density of high-speed signals, proper power/ground network design is required for adequate noise isolation. The power methodology encompasses all circuit elements from transistor layout and appropriate deep N-well encapsulation through sub-block power partitioning all the way to the package power pins. Appropriate high speed and power pin placement is a key to successful customer board layout designs.

E. Layout integration

Two of the main issues with layout integration were:

i) **Power isolation:** Each unique circuit block in any high-speed channel has its own power and ground network, coupling capacitance and bumps. A deep N-well layer was used to isolate sensitive layout blocks and reduce noise interference.

ii) **FPGA Logic Array - Hard IP Integration:** The layouts for these two areas were drawn in slightly different design rules, though derived from the same fabrication (TSMC) process. Full chip layout verification presented a knotty problem. A simple solution was adopted – isolate the two areas with a ring around the Hard IP of sufficient width. Only metal routed signals traversed this ring. The Hard IP and FPGA Core were independently verified to be DRC clean and were simply merged into a single database.

Testing

Different component configurations, testing patterns and equipment were used to verify functionality of the silicon. One such setup is shown in Fig. 8.




Fig. 8 Characterization Setup

Results: HSSI

With 75% LE usage, all 4 HSSI quads (20 channels) are operating at least at 3.24Gbps under XAUI external serial loopback test.

Fig. 9 shows an example of a transmit eye diagram (no pre-emphasis). The silicon measurement matches

simulation within a few percent. Fig. 10 shows the eye diagram at the far-end after 40' FR4 interconnect using FR4, two high speed connectors, four SMA connectors, 6ft of cable and without pre-emphasis. Simulation and Silicon data are in very close agreement. Pre-emphasis on TX driver and receiver equalization can be dynamically adjusted for optimum link performance.




Fig. 9 Near End Correlation




Fig. 10 Far End Correlation

The DPA also successfully met the 1Gbps target.


Acknowledgements

The author would like to thank the members of the Altera development teams – Layout, CAD, Product Engineering, Software, Applications, and Product Planning – whose valuable contributions are greatly appreciated.

References

- (1) S.C.Wong et al., "CMOS Erasable Programmable Logic Device with Zero Standby Power", ISSCC Digest of Technical Papers, Feb. 1986, PP 242-243
- (2) M.J.Allen, "A 6 nanosecond CMOS EPLD with mW Standby Power", CICC '89, May 15-18, 1989
- (3) D. Lewis et al., "The Stratix™ Routing and Logic Architecture", in FPGA'03, February 23-25, 2003
- (4) Altera Product Data Sheet
- (5) IEEE Draft P802.3ae/D5.0

Layout Snapshot





101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
www.altera.com
Applications Hotline:
(800) 800-EPLD
Literature Services:
literature@altera.com

Copyright © 2005 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

All copyrights reserved.