Implementation of Wave-Pipelined Interconnects in FPGAs

Terrence Mak[1]*, Crescenzo D'Alessandro[2], Pete Sedcole[1], Peter Y.K. Cheung[1],
Alex Yakovlev[2] and Wayne Luk[3]
[1]Department of Electrical and Electronic Engineering, [3]Department of Computing,
Imperial College London, London, UK
[2]School of Electrical, Electronic and Computer Engineering,
Newcastle University, UK

Abstract

Global interconnection and communication at high clock frequencies are becoming more problematic in FPGA. In this paper, we address this problem by presenting an interconnect wave-pipelining strategy, which utilizes the existing programmable interconnects fabrics to provide high-throughput communication in FPGA. Two design approaches for interconnect wave-pipelining, using simple clock phase shifting and asynchronous phase encoding, are presented in this paper. Experimental results from a Xilinx Virtex-5 FPGA device are also presented.

1 Introduction

Recently, several novel designs of global communication link have been proposed. These new proposals provide an energy efficient and high throughput alternative to the conventional point-to-point interconnections. Notably, interconnect wave-pipelining [4, 3, 2] was introduced as an effective solution to increase the global interconnection throughput. It offers an opportunity to overcome the ever-increasing global interconnection delay problems and can potentially be adopted in FPGAs.

This paper presents two different approaches to realize wave-pipelined interconnects in FPGAs. The first approach is a synchronous design using clock phase shifting at the receiver end to sample the wave-pipelined data. The second approach is using the asynchronous phase-encoding technique to encode data with differential signaling in two wires. This approach is robust due to the differential signaling and can potentially provide higher throughput. We implemented these two approaches with wave-pipelining in a state-of-the-art FPGA and hardware testing results are reported.

*Email: t.mak@imperial.ac.uk; T. Mak is gratefully acknowledge supports from the Croucher Foundation.

2 FPGA Implementation

2.1 Simple Clock Phase Shifting

Clock shifting approach is a simple and efficient approach to implement wave-pipelining. It requires calibration of phase-shifting for the receiver clock after the placement and routing of the link. Sender and receiver blocks are clocked by snd_clk and rsv_clk respectively. The receiver clock has a phase relationship ϕ with respect to the sender clock. The testing circuit also comprises of a test pattern generator at the sender and an identical pattern generator at the receiver, which will be used to verify the incoming data. The error results will be registered and counted by the Microblaze processor in the FPGA. For a Xilinx FPGA, clock phase shifting can be realized by using the Digital Clock Manager (DCM) embedded module, which can provide reasonably high resolution phase shifting of the clock. Alternatively, the phase can be controlled by inserting delay logics, such as clock buffers.

The advantage of this approach is that it does not require any extra dedicated logic in order to implement the wave-pipelining. However, the design requires proper calibration of the phase. Also, FPGA without embedded phase lock-loop (PLL) or DCM would be difficult to provide exact matching of the clock phase at the receiver end.

2.2 Asynchronous Phase Encoding

Phase encoding [1] is an asynchronous signalling approach that the clock is embedded in the encoded data. The concept is employing the order of events on a pair of wires to indicate the bit value. The scheme allows the use of both rising and falling edges for transmission, allowing a natural multiplexing of two channels onto the same link. This

978-0-7695-3098-7/08 $25.00 © 2008 IEEE
DOI 10.1109/NOCS.2008.32

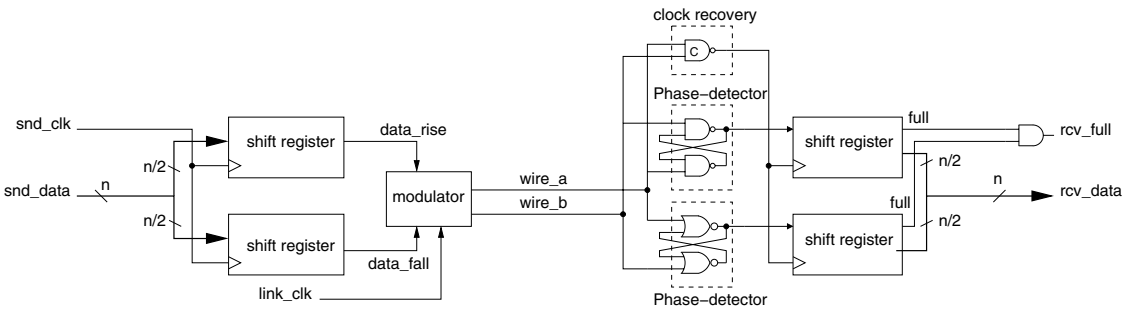

Figure 1. Schematic for a dual-channel phase-encoding link.

can be exploited to increase the bitrate of the link, while maintaining a reduced operating frequency. Fig. 1 shows the schematic diagram of a dual-channel link with phase encoding. The time separation between the edges is immaterial and only needs to satisfy the setup/hold condition of the receiver's phase detector in order to ensure correct operation.

The modulator circuit provides two sets of data inputs for the data at the rising and falling edge of the reference clock (link_clk). The circuit creates a phase difference at wire_a and wire_b based on the propagation gates delay. Matching of the gates delay in the circuit is important to maintain a reasonable phase encoding. In FPGA, this can be achieved by applying a tight area constraint for the placement to make sure that interconnect delay in the circuit is minimized. Furthermore, delay elements (τ) can be inserted to increase the phase difference and, especially, this can avoid local interconnect delay dominant the delay path in the design.

The advantages of phase-encoding as a serialized signalling scheme for FPGA can be summarized as follows. Firstly, The scheme can be readily implemented in FPGA using only reconfigurable logics. Secondly, the modulator will regenerate the absolute phase relationship between edges, within the limits of wave-pipelining i.e. without the need for clocked latches. Thirdly, it is robust to single-event upsets, as a single event on one wire because of the differential signalling. However, an important limitation to the use of phase-encoded links consists in the presence of logic feedback loops, necessary to the design of reliable phase-encoding. The link can be designed without feedback loops if delay lines are used, but this would reduce the overall robustness, introducing opportunities for process variations-induced faults.

3 Results

The two designs of interconnect wave-pipelining has been implemented in a Xilinx Virtex-5 XC5VLX50 FPGA device (with speed grade -1). For a 75-tile length interconnection with 5.5 ns propagation latency, the maximum frequency of data rate is at around 185 MHz. With phase shifting, there will be more than one data bit traversing the line simultaneously. For frequency at around 350 MHz, which is almost doubles the original maximum frequency. For the phase encoding design, it can achieve a maximum frequency 170 MHz of the link_clk at Xilinx Virtex-5 XC5VLX50 FPGA. Since the link transmit data with both rising and falling edges, the link can achieve a transmission rates of 340 Mb/s. For a link with length 75 tiles and with transmission rate 340 Mb/s, it doubles the synchronous transmission rate with wave-pipelining. Incorporation of wave-pipelining design into real applications and aiming to reduce area and power will be our future work.

References

[1] C. D'Alessandro, D. Shang, A. Bystrov, and A. Yakovlev. PSK Signalling on SoC Buses. In *Proceedings of PATMOS*, 2005.

[2] R. R. Dobkin, Y. Perelman, T. Liran, R. Ginosar, and A. Kolodny. High rate wave-pipelined asynchronous on-chip bit-serial data link. In *Proceedings of ASYNC*, 2007.

[3] S.-J. Lee, K. Kim, H. Kim, N. Cho, and H.-J. Yoo. Adaptive network-on-chip with wave-front train serialization scheme. In *Proceedings of Symposium on VLSI Circuits Digest of Technical Papers*, 2005.

[4] J. Xu and W. Wolf. A Wave-Pipelined On-chip Interconnect Structure for Networks-on-Chips. In *Proceedings of the 11th Symposium on High Performance Interconnects*, 2003.