AFRL-RI-RS-TR-2015-047

EXPERIMENTAL 3D ASYNCHRONOUS FIELD PROGRAMMABLE GATE ARRAY (FPGA)

CORNELL UNIVERSITY

MARCH 2015

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88th ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2015-047 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/ S /

THOMAS E. RENZ
Work Unit Manager

/ S /

MARK H. LINDERMAN
Technical Advisor, Computing
& Communications Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) MARCH 2015		2. REPORT TYPE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) OCT 2011 – OCT 2014	
4. TITLE AND SUBTITLE EXPERIMENTAL 3D ASYNCHRONOUS FIELD PROGRAMMABLE GATE ARRAY (FPGA)				5a. CONTRACT NUMBER FA8750-12-2-0035	
				5b. GRANT NUMBER N/A	
				5c. PROGRAM ELEMENT NUMBER 62788F	
6. AUTHOR(S) Rajit Manohar				5d. PROJECT NUMBER T2PM	
				5e. TASK NUMBER CO	
				5f. WORK UNIT NUMBER RN	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Cornell University Office Of Sponsored Programs 373 Pine Tree Rd Ithaca NY 14850-2820				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RITB 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI	
				11. SPONSOR/MONITOR'S REPORT NUMBER AFRL-RI-RS-TR-2015-047	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. PA# 88ABW-2015-1039 Date Cleared: 11 MAR 2015					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT 3D technology has the potential to significantly enhance the capabilities of a computing platform through tight integration of multiple levels of logic through low-cost vertical interconnects. The goal of this effort is to evaluate the potential of 3D technology in the context of an asynchronous Field Programmable Gate Array developed for the Air Force Research Laboratory through an experimental demonstration in a 3D technology in collaboration with Albany's College of Nanoscale Science and Engineering.					
15. SUBJECT TERMS 3D Technology, vertical interconnects, AFPGA, FPGA					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT U	18. NUMBER OF PAGES 22	19a. NAME OF RESPONSIBLE PERSON THOMAS E. RENZ
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) NA

Table of Contents

List of Figures.....	ii
1. Summary.....	1
2. Introduction	2
3. Methods and Procedures.....	3
3.1. Asynchronous FPGA Overview	3
3.2. 3D Routing with Through Silicon Vias.....	4
3.3. Alternate Signaling Approaches for 3D	8
4. Results and Discussion.....	11
4.1. Study of 3D Communication.....	11
4.2. 3D FPGA Design.....	13
5. Conclusions	14
6. References	15
7. Acronyms	17

List of Figures

Figure 1. 2D FPGA Architecture With Switch Box (SB), Logic Box (LB), and Logic Routing (LR).....	3
Figure 2. Layout for Switch Point in the AFPGA.	5
Figure 3. Sample Test Structure on 3D Run.	6
Figure 4. Scale of AFPGA Tile versus TSVs, Assuming 25 μ m TSV Dimensions.....	7
Figure 5. Design for 3D With Denser TSVs and Cascaded 2D/3D Buffers.....	8
Figure 6. Decoding the Vdd/2 State Using Level Shifters.....	9
Figure 7. Different Circuit Options for Vdd/2 Driver.....	10
Figure 8. Comparing Different Circuit Styles and Signaling Schemes for Throughput and Energy.	12
Figure 9. Unified Pareto Optimal Frontier for Different Technologies.....	13

1. Summary

Previous research on asynchronous FPGA architectures at Cornell resulted in the development of a new high performance reconfigurable fabric. This funded effort evaluates the potential of 3D integration to impact FPGA architectures, and more generally quantifies the communication costs of 3D vs 2D integration through fabrication experiments in collaboration with Prof. Geer's group at SUNY CNSE. New techniques for asynchronous communication that minimize the use of wires were developed that are superior to conventional approaches when compared on bandwidth density and energy. A design was submitted to CNSE for fabrication in their project.

2. Introduction

Conventional computer architectures such as uniprocessor, multicore, or massively parallel processors have a number of limitations when it comes to the performance requirements of DoD missions. These limitations include high power consumption during floating-point computations, high latency of global reduction operations causing performance degradation of parallel simulations, and high latency and energy cost of retrieving data from off-chip memories. Cornell has been collaborating with the Air Force Research laboratory (AFRL) on a new architecture that combines an energy-efficient embedded processor architecture with a high-performance asynchronous Field Programmable Gate Array (FPGA) to address some of these requirements.

The overhead of a programmable interconnect in an FPGA architecture is significant. The vast majority (between 70%-90% depending on the architecture) of the area, energy, and delay of an FPGA fabric is in its flexible interconnect network. Any improvements to interconnect technology could have a major impact on the efficiency of FPGAs. Hence, 3D integration is a technology that has the potential to significantly enhance an FPGA architecture. By stacking multiple device layers, it is possible to create a dense computation fabric where the third dimension reduces path lengths and the energy cost of moving data from one part of the computation to another. A natural way to explore this potential is to extend the reconfigurable fabric to the third dimension, allowing 3D connectivity to be under software control.

This report contains a summary of the work that was conducted in evaluating the potential of 3D technology in the context of an existing FPGA developed in collaboration with AFRL.

3. Methods and Procedures

The Cornell Asynchronous Very Large Scale Integration (VLSI) and Architecture group previously developed a high-performance FPGA fabric for general-purpose computing. Compared to the state-of-the-art commercial FPGAs from industry, the performance of the fabric was three times higher—a significant improvement. Compared to the best previously developed asynchronous FPGAs, the Cornell FPGA was almost twenty times faster in terms of application throughput [1, 2, 3]. This dramatic performance increase makes the fabric ideally suited to be integrated into a system containing a high-performance microprocessor.

3.1. Asynchronous FPGA Overview

In terms of the major building blocks, the asynchronous FPGA (AFPGA) architecture looks like a traditional synchronous island-style FPGA such as a Xilinx Virtex [4]. The FPGA contains a configurable logic block (LB) and a configurable interconnect, with the interconnect being broken down into global block-to-block connectivity (global switch box routing, or SB) and connectivity internal to each logic block (logic routing, or LR). Figure 1 shows a high-level view of a generic modern FPGA architecture.

Figure 1. 2D FPGA Architecture With Switch Box (SB), Logic Box (LB), and Logic Routing (LR).

The major differentiating feature of the AFPGA versus a conventional FPGA architecture is the underlying computation model used to implement the configurable fabric. Instead of thinking of computation in terms of gates and registers, the AFPGA implements a computation specified by a dataflow graph [5]. In the dataflow graph model, computation is described by operations on data values or “tokens” flowing through the graph. Tokens correspond to valid data items being processed by elements of the dataflow graph. Nodes in the dataflow graph include function blocks that can perform computation, as well as routing elements for sending tokens to the appropriate destinations. Token arrival at a dataflow node can be thought of as an “event” that triggers activity in the AFPGA.

The key performance amplifier in the AFPGA is its flexible routing network. A conventional FPGA has over 70% of its delay in the routing network [6]. Since the AFPGA operates using a dataflow model, pipeline stages corresponding to queues can be introduced into the routing network without impacting the correctness of the computation being performed by the AFPGA! This means that designs can benefit from pipelining without the additional cost required from electronic design automation (EDA) tools to support interconnect pipelining. In the first AFPGA implementation, pipelined stages were introduced in the switch boxes in the AFPGA interconnect [1,2].

The nature of the pipelined interconnect makes the entire AFPGA highly modular. In particular, because communication between components on the AFPGA uses the dataflow model, the delay of the communication link is not part of the interface specification. This enables a highly modular approach to the design of the AFPGA, where sub-blocks can be pre-placed without significantly impacting performance. Indeed, if data flow between one sub-block and another is unidirectional (as in a computation pipeline), there is no loss in throughput by using a modular approach to synthesis and place-and-route.

The impact of aggressive pipelining on the overall performance of the AFPGA is significant. In a 0.18 μ m feature size, the measured peak performance of the AFPGA architecture was 674 MHz. For reference, the baseline Xilinx architecture in a similar feature size performs at 240MHz [3]. More important, first pass synthesis results for a variety of benchmarks demonstrate robust performance. For example, a synthesized Finite Impulse Response (FIR) filter core would exhibit a performance of 75% of the peak performance of the AFPGA.

3.2. 3D Routing with Through Silicon Vias

Through silicon vias (TSVs) are a promising new technology that permit vertical interconnects between multiple device layers. The approach extends the normal vias that connect multiple levels of metal to connections between multiple wafers or chips. TSVs permit multiple silicon wafers (or chips) to be vertically stacked and then interconnected in a dense fashion. The net effect is a chip stack that has active devices integrated with both planar and vertical connectivity—a 3D integrated system.

The third dimension can have the potential to reduce wiring costs. At an abstract level, one can pack $O(R^3)$ instead of $O(R^2)$ densely interconnected devices, where R is the physical diameter of the system. This should lead to a reduction in wiring costs for a fixed number of devices—from $O(N^{1/2})$ to $O(N^{1/3})$, where N is the number of devices in the system. 3D chip stacking has been proposed as a way to improve microprocessor performance [7,8], complex systems-on-chip [9], as well as FPGA designs [10,11,12,13].

There are a variety of approaches to manufacturing TSVs and ensuring high yield, but the net effect of these approaches on design is captured by a set of design rules for TSV layout, similar to design rules for other physical geometry in VLSI. These design rules have a significant impact on the way TSVs can be integrated into a design. For this project, we worked in collaboration with Prof. Geer's group at SUNY's College for Nanoscale Science and Engineering (CNSE), whose group provided the manufacturing expertise for 3D TSVs in IBM's 65nm CMOS process. In the discussion that follows, we refer to this specific approach to TSV manufacturing.

The goal of the first 3D design was two-fold: for us to get some design experience with the TSV technology, and for Prof. Geer's group to evaluate the density of TSVs. Since the application driver for this project is an AFPGA, a highly replicated component of the AFPGA was selected as a test structure for the first 3D design. This structure contains:

1. Pure 2D routing. This corresponds to the baseline 2D design that we will compare against. This also provides a mechanism to evaluate the impact of 3D processing on 2D performance.
2. Pure 3D switch point routing. This test provides an evaluation of vertical routing, and the 2D overhead of the 3D switch point architecture
3. Mixed 2D and 3D switch points. This test provides an evaluation of the routing as envisioned in the final AFPGA. The mixture of switch points is needed because the TSV pitch limits the number of 3D paths supported per AFPGA tile.

The basic building block in the global routing fabric of an AFPGA is the pipelined switch point. A standard pipelined switch-point supports 2D routing by having input and output communication links in the 2D plane, as well as configuration information that specifies how these links are connected to each other. The switch point is pipelined, which means that there is an asynchronous pipeline buffer embedded in the switch point. This permits high-throughput operation, because electrical signaling paths are always kept short due to internal pipelining within the switch point. Figure 2 shows the physical layout for a dense switch point in an existing 2D AFPGA architecture

Figure 2. Layout for Switch Point in the AFPGA.

The switch point has two buffers, because a pair of horizontal and vertical communication links can support up to two independent parallel connections that require separate physical buffering.

Figure 3 shows the typical test structure for the first 3D run. A chain of switch points

Figure 3. Sample Test Structure on 3D Run.

(either pure 2D, pure 3D, or mixed) is constructed with a source generating data and a sink consuming data. An internal signal is probed, and connected to a frequency divider for low frequency external measurement. Any circuit that limits the total throughput will limit the frequency reported by the divider, because all data values travel through all the circuits in the test structure.

TSV design rules imply that the pitch between adjacent TSVs is on the order of tens of microns. A distance of 10 μm in a 65nm manufacturing technology is a distance that is about 308λ , where λ is half the feature size. Standard wiring pitches are below 10λ —more than 30x lower than practical manufacturing limits for TSVs today. Hence, integration of TSV interconnects cannot be done in a manner similar to conventional wiring.

Figure 4. Scale of AFPGA Tile versus TSVs, Assuming 25 μm TSV Dimensions.

The impact of the physical size of TSVs versus the physical size of an AFPGA tile that contains logic and routing is illustrated in Figure 4. Since known reliable TSV manufacturing sizes are in the 25 μm regime (which means that the TSVs themselves are 25 μm by 25 μm , with a pitch of 50 μm between adjacent TSVs), Figure 4 shows the size of known reliable TSVs juxtaposed with a single AFPGA tile. The “long edge” of the AFPGA tile is 160 μm , which means that we can fit three TSVs per tile. Contrast this with the number of horizontal and vertical interconnect wires that cross a tile—192 in the current design. Hence, tight integration at the density of the 2D interconnect in an AFPGA is not possible given current TSV density. Therefore, we examined an alternate approach where we provide some limited vertical connectivity that is *shared* by the 2D interconnect.

Progress in TSV size by Prof. Geer's group as part of this project enabled significantly improved TSV density, enabling us to use twelve TSVs per edge. To support this extra connectivity, additional configuration memory is required as part of the AFPGA state. Instead of modifying the 2D connectivity, our design adds a *cascaded* 3D switch point that is shared by multiple 2D tracks. This design is illustrated in Figure 5, showing the original 2D AFPGA tile enhanced with 3D TSVs (circles), additional buffering, and 3D configuration memory to control vertical connectivity.

Figure 5. Design for 3D With Denser TSVs and Cascaded 2D/3D Buffers.

3.3. Alternate Signaling Approaches for 3D

There is a second option that can be used to improve vertical connectivity that was studied as part of this project. In an asynchronous communication scheme, the signaling wires must convey not just *what* data is being communicated, but also *when* data is being communicated [14]. In other words, a communication scheme to send a single bit requires at least three states: sending data 0, sending data 1, and not sending any data. The most common mechanism to do this encoding is to use two wires per bit, and three states of those two wires: 00 representing no data being communicated; 01 representing data 0 being communicated; and 10 representing data 1 being communicated. Finally, because there is no global clock, a receiver must indicate that data transmitted has been received successfully. This indication is provided by a third *acknowledge* wire. Hence, a single, independently communicated bit uses three wires for signaling [15]. This is part of why Figure 5 shows three TSVs grouped per buffer.

Instead, other more complex signaling schemes are possible, and we studied two other approaches as part of this project.

- Single track full buffering (STFB). In this scheme, two wires are used to communicate a bit, but the wires are *bidirectional* [16,17]. The state 00 indicates no data is being transmitted. To transmit data 0, the sender changes the states of the wires to 01. To transmit data 1, the sender changes the states of the wires to 10. This is similar to conventional signaling so far. However, to *acknowledge* that data has been received, the *receiver* restores the state of the wires to 00.
- Asynchronous ternary logic signaling (ATLS). In this scheme, two wires are also used to communicate a bit [18]. Data is encoded on a single wire by using a third voltage level $V_{dd}/2$, where V_{dd} is the nominal supply voltage. The acknowledge wire remains unchanged compared to the conventional signaling approach. To send a data 0, the data wire is lowered from $V_{dd}/2$ to 0V. To send a data 1, the data wire is raised from $V_{dd}/2$ to V_{dd} .

Single track asynchronous ternary signaling (STATS). This new approach developed for this project combines the two previously described schemes, using one wire to communicate a bit where the wire is used bidirectionally [19]. Initially the wire is at voltage $V_{dd}/2$. To transmit a 0, the sender lowers the voltage of the wire to 0V. To transmit a 1, the sender raises the voltage of the wire to V_{dd} . Once the data has been received, the receiver resets the voltage of the wire to $V_{dd}/2$, acknowledging receipt of data.

Figure 6. Decoding the $V_{dd}/2$ State Using Level Shifters.

Using $V_{dd}/2$ as an additional voltage reduces noise margins, but given the area overhead of current generation TSVs, we believed this was an acceptable trade-off to triple the density of vertical interconnects possible compared to more traditional

asynchronous signaling schemes. All three reduced wire signaling approaches were evaluated for this project.

The STATS link requires circuits that can set a wire to $V_{dd}/2$, as well as circuits to detect when a signal reaches the $V_{dd}/2$ threshold. To decode the ternary voltage levels, a pair of level shifters is used as illustrated in Figure 6. Note that one of the level shifters converts voltages on a wire in the range $(0, V_{dd}/2)$ to $(0, V_{dd})$, while the other converts the range $(V_{dd}/2, V_{dd})$ to $(0, V_{dd})$. This combination creates two different logic signals whose combination can be used to determine whether the voltage on the wire is 0, $V_{dd}/2$, or V_{dd} .

To *drive* the wire to $V_{dd}/2$, we evaluated three different schemes. The pass gate scheme drives the link by connecting it to an external $V_{dd}/2$ supply. It is the most conservative scheme, and also the slowest. The second scheme is the *self-invalidating driver*, where a more traditional driver is used along with a $V_{dd}/2$ detection circuit that triggers a self shut-off of the driver. This is aggressive, and very sensitive to a number of circuit parameters such as the delay of the $V_{dd}/2$ detection circuit and the load being driven by the circuit. The third scheme is the *shorted inverter* approach, which exploits the voltage transfer characteristic of a CMOS inverter. In this approach, the wire is driven by a shorted inverter circuit—an inverter whose output is connected to its input. While this scheme is fast, it uses significant energy because it temporarily shorts V_{dd} to GND. These three options are shown in Figure 7.

Figure 7. Different Circuit Options for $V_{dd}/2$ Driver.

4. Results and Discussion

4.1. Study of 3D Communication

The design for different 3D communication structures was submitted to Albany CNSE for a fabrication run which didn't occur until after this project was scheduled for completion. We are awaiting test results to quantify the benefits of 3D vs 2D connections.

We performed detailed analog simulations using Synopsys' HSPICE simulation package to quantify the benefits of different approaches to asynchronous communication in the context of 3D integration. To simulate TSV links between device tiers, we used public information about the electrical properties of TSV links and ensured those were relatively consistent with the TSV links being fabricated by SUNY CNSE. Simulations also accounted for electrical coupling between neighboring vertical TSVs. In the results, we report throughput *per TSV*, since TSVs by themselves are a very scarce resource. We also examined different technology nodes while keeping the TSV dimensions relatively fixed because CMOS technology scaling and TSV scaling are not coupled.

For three wire signaling, we also looked at a number of different circuit topologies possible for asynchronous signaling, including a standard weak-conditioned half buffer circuit (WCHB), as well as a more aggressive timed circuit (RQDI, for relaxed quasi delay insensitive) that uses two-phase communication protocols rather than the conventional return-to-zero four phase communication protocols [18].

Figure 8 contains a summary of simulation results. The axes show energy per data token versus throughput per TSV. We used an optimization package to vary circuit parameters such as device size, circuit topology, number of repeaters, etc. to determine feasible points in the energy-throughput space for each circuit family. Each curve in Figure 8 corresponds to the Pareto-optimal frontier for the given circuit family. In this analysis, the best circuit choices vary depending on the required throughput—at low frequencies, the WCHB style provides good throughput at the lowest energy per data token. At the highest frequencies, only the STATS approach is feasible. We can combine all these Pareto-optimal curves into a single unified frontier of feasible points.

Figure 8. Comparing Different Circuit Styles and Signaling Schemes for Throughput and Energy.

Figure 9 shows the unified Pareto-optimal curves for three different technologies: 90nm, 65nm, and 45nm. While we originally started with five different circuit and signaling options, the ATLS option was found to be inferior to one of the other four circuit families for all possible points in the design space.

In 90nm and 65nm, low frequency designs use the WCHB circuit style. As the frequency requirement is increased, the best circuit choice switches to RQDI and then STFB in terms of providing the lowest energy per data token. At the highest frequency, the STATS style dominates.

As technology scales, the STATS circuit style proves to be the best choice across a wider range of throughputs. This is to be expected, because it uses at least 50% fewer TSVs than any other design style, and the overhead due to the extra voltage supply and $V_{dd}/2$ detection and driver circuit becomes smaller relative to the cost of the TSV itself.

Figure 9. Unified Pareto Optimal Frontier for Different Technologies.

4.2. 3D FPGA Design

A three-tier 16x16 clustered AFPGA was submitted for fabrication in collaboration with SUNY CNSE with a total capacity of 3,072 four-input look-up tables (LUTs) along with hardware support for arithmetic as well as dynamic reconfiguration. We are currently awaiting results from fabrication, as the timeline for fabrication exceeds the end of this project due to additional post-processing time required for 3D manufacturing.

5. Conclusions

Previous research on asynchronous FPGA architectures at Cornell resulted in the development of a new high performance reconfigurable fabric. This funded effort evaluates the potential of 3D integration to impact FPGA architectures, and more generally quantifies the communication costs of 3D vs 2D integration through fabrication experiments in collaboration with Prof. Geer's group at SUNY CNSE. New techniques for asynchronous communication that minimize the use of wires were developed that are superior to conventional approaches when compared on bandwidth density and energy.

A successful design was delivered to the fabrications group at CNSE. The design passed all design rule tests. The fabrication won't take place until after the end of this effort, so testing results are not included in this report.

6. References

- [1] John Teifel and Rajit Manohar, “Programmable Asynchronous Pipeline Arrays”, *Proceedings of the 13th International Conference on Field Programmable Logic and Applications*, Lisbon, Portugal, September 2003.
- [2] John Teifel and Rajit Manohar, “Highly Pipelined Asynchronous FPGAs”, *12th ACM International Symposium on Field-Programmable Gate Arrays*, Monterey, CA, February 2004.
- [3] David Fang, John Teifel, and Rajit Manohar, “A High-Performance Asynchronous FPGA: Test Result”, *2005 IEEE Symposium on Field-Programmable Custom Computing Machines*, April 2005.
- [4] Xilinx. VirtexTM 2.5V Field Programmable GateSarrays. Xilinx Data Sheet, 2002.
- [5] J. B. Dennis, “The Evolution of ‘Static’ Data-Flow Architecture”, *Advanced Topics in Data-Flow Computing*, J.-L. Gaudiot and L. Bic, editors, Prentice-Hall, 1991.
- [6] I. Kuon, R. Tessier, J. Rose, “FPGA Architectures: Survey and Challenges”, *Foundations and Trends in Electronic Design Automation*, Vol. 2, No. 2, pp. 135-253, 2007.
- [7] W. Rhett Davis, John Wilson, Stephen Mick, Jian Xu, Hao Hua, Christopher Mineo, Ambarish M. Sule, Michael Steer, and Paul D. Franzon, “Demystifying 3D ICs: The Pros and Cons of Going Vertical”, *IEEE Design and Test*, Vol. 22, No. 6, pp. 498–510, 2005.
- [8] Christianto C. Liu, Ilya Ganusov, Martin Burtscher, Sandip Tiwari, “Bridging the Processor-Memory Performance Gap With 3D IC Technology”, *IEEE Design and Test of Computers*, 2005.
- [9] Christianto C. Liu, Jeng-Huei Chen, Rajit Manohar, and Sandip Tiwari, “Mapping Multimedia Applications to 3-D System-on-Chip”, *Proceedings of the 2005 IEEE International Symposium on Circuits and Systems*, May 2005.
- [10] David Fang, Christopher LaFrieda, Song Peng, and Rajit Manohar, “A 3-Tier Asynchronous FPGA”, *Proceedings of the 23rd International VLSI/ULSI Multilevel Interconnection Conference (VMIC)*, September 2006.
- [11] Cristinel Ababei, Pongstorn Maidee, and Kia Bazargan, “Exploring Potential Benefits of 3D FPGA Integration”, *14th International Conference on Field Programmable Logic and Applications (FPL 2004)*, pp. 874–880, 2004.
- [12] Mingjie Lin, Abbas El Gamal, Yi-Chang Lu, and Simon Wong, “Performance Benefits of Monolithically Stacked 3D-FPGA”, *Proc. Internation Symposium on Field Programmable Gate Arrays (FPGA ’06)*, pp. 113–122, 2006.

- [13] Arifur Rahman, Shamik Das, Anantha P. Chandrakasan, and Rafael Reif, “Wiring Requirement and Three-Dimensional Integration Technology for Field Programmable Gate Arrays”, *IEEE Trans. Very Large Scale Integr. Syst.*, Vol. 11, No. 1, pp. 44–54, 2003.
- [14] Rajit Manohar, “A Case for Asynchronous Computer Architecture”, *Proceedings of the ISCA Workshop on Complexity-Effective Design*, June 2000.
- [15] J. Brzozowski and C.-J. Seger, **Asynchronous Circuits**, Springer-Verlag, 1995.
- [16] K. van Berkel and A. Bink, “Single-Track Handshake Signaling with Applications to Micropipelines and Handshake Circuits”, *Proc. 2nd International Symposium on Asynchronous Circuits and Systems*, 1996.
- [17] M. Nystroem, **Asynchronous Pulse Logic**, Ph.D. thesis, California Institute of Technology, 2002.
- [18] T. Felicijan and S. Furber, “An Asynchronous Ternary Logic Signaling System,” *IEEE VLSI*, Vol. 11, No. 6, pp. 1114–1119, 2003.
- [19] Jonathan Tse, Benjamin Hill, and Rajit Manohar, “A Bit of Analysis on Self-Timed Single-Bit On-Chip Links”, *Proceedings of the 19th IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)*, May 2013.

7. Acronyms

Acronym	Expanded Form
2D	Two dimensional
3D	Three dimensional
AFPGA	Asynchronous Field Programmable Gate Array
AFRL	Air Force Research Laboratory
ATLS	Asynchronous ternary logic signaling
CMOS	Complementary Metal Oxide Semiconductor
CNSE	College of Nanoscale Science and Engineering
EDA	Electronic Design Automation
FIR	Finite impulse response
FPGA	Field Programmable Gate Array
IBM	International Business Machines
LB	Logic block
LR	Logic routing
LUT	Lookup Table
RQDI	Relaxed quasi delay insensitive
SB	Switch box
SRAM	Static Random Access Memory
STFB	Single track full buffer
STATS	Single track asynchronous ternary signaling
SUNY	State University of New York
TSV	Through silicon via
VLSI	Very Large Scale Integration
WCHB	Weak conditioned half buffer