A High-Throughput Low-Power Fully Parallel 1024-bit ¹/₂-Rate Low Density Parity Check Code Decoder in 3-Dimensional Integrated Circuits*

Lili Zhou, Cherry Wakayama, Nuttorn Jangkrajarng, Bo Hu, and C.-J. Richard Shi

Department of Electrical Engineering, University of Washington, Seattle, WA 98195, USA

Abstract - A 1024-bit, $\frac{1}{2}$ -rate fully parallel low-density parity-check (LDPC) code decoder has been designed and implemented using a three-dimensional (3D) 0.18µm fully depleted silicon-on-insulator (FDSOI) CMOS technology based on wafer bonding. The taped-out 3D decoder with about 8M transistors was simulated to have a high throughput of 2Gb/s and a low power consumption of only 430mW using 6.4µm by 6.3µm of die area. The 3D implementation is estimated to offer more than 10x power-delay-area product improvement over its corresponding 2D implementation. This first large-scale 3D ASIC with fine-grain (5µm) vertical interconnects is made possible by jointly developing a complete automated 3D design flow from a commercial 2-D design flow combined with the needed 3D-design point tools.

I. Introduction

Low density parity check (LDPC) codes are emerging as standard methods of channel encoding and error correcting for many wireless standards, due to their near Shannon-limit error correction performance [1] and the progress in semiconductor fabrication technologies that allow very large scale integration of circuit functionality. The LDPC block-parallel message passing decoding algorithm and its fully-parallel implementation architecture yield the high-throughput error-correction capacity necessary for large-volume communication and data storage applications. However, the implementation leads to the following interconnect design challenges [2]:

(1) The average wire length can be 3mm for a 7.5mm x7.0 mm die implementation (half of the die size).Therefore, specific CAD tools have been developed.

(2) The wiring takes more silicon area and power dissipation than the logic itself.

(3) Only 50% area utilization for logic was achieved due to routing congestion.



Figure 1: Cross-section of 3-tier 3D integration.

To address this interconnect design challenge, we explore the use of a 3-dimensional integrated circuit process.

More specifically, we use MIT Lincoln Lab's wafer bonding 3D process, which stacks 3 wafers each composing of a single layer of transistors with 3 layers of metal wires (called one-tier) formed on fully-depleted silicon-on-insulator (FDSOI) substrates [3]. Figure 1 shows the cross-section view of the 3-tier 3-D IC integration.

We note that all the previous 3D IC designs are limited either to simple logic devices, or to circuits of regular structures such as photo sensors, memories and field-programmable gate arrays (FPGAs). This is primarily due to the lack of 3-D CAD tools and a complete 3D design flow to handle the ASIC design complexity. To accomplish this challenging LDPC ASIC design in 3D, we have developed the needed 3D-CAD tools to automate the 3D IC implementation process. A complete 3D design methodology and design flow is developed based on a commercial 2D design flow augmented with the needed 3D design tools.

II. Fully Parallel LDPC Decoder Architecture

Our 3D LDPC decoder is based on the classical message-passing (belief-propagation) algorithm, which maps extremely well to a parallel decoder architecture that can be represented by a bipartite graph, directly instantiated to hardware [2]. As illustrated in Fig. 2, there are two types of computing nodes (*variable* nodes and *check* nodes) that perform all the logical calculations, edges representing the required interconnect as defined by the very sparse parity check matrix. A 1024-bit ½-rate code decoder requires 1024 variable nodes and 512 check nodes.



Figure 2. The Tanner graph of a LDPC code.

The main data path is designed as 16 parallel three-stage pipelines (Fig. 3); this allows the decoder to achieve the high throughput of 2Gb/s with the clock frequency at 128MHz (128MHz*16=2Gb/s).



Figure 3. Fully parallel 3-stage pipelined LDPC architecture.

^{*} This research was sponsored by the U.S. DARPA 3D-IC Program under Grant No. N66001-05-1-8918 monitored by Navy SPAWAR, San Diego.

III. 3D LDPC Design

Our 3D design methodology is based on partitioning a 3D-design into 2-D designs at the fine-grain level. An in-house 3D placement tool has been developed that places computing nodes on all the 3 tiers with the objectives to minimize the area, routing density, total wire length and 3D-vias. We have also developed in-house programs for 3D routing, buffer insertion, and circuit-verse-schematic (CVS) checking. Figure 4 shows the 3D LDPC design flow.

IV. Results and 2-D Comparison

The 3D LDPC decoder has been taped out based on the 3D MIT Lincoln Lab FD-SOI process. It contains about 8M transistors using $6.4\mu m$ by $6.3\mu m$ of 3D die area. The top view of 3-tier LDPC layout is shown in Figure 5, where over 10,000 dense 3D vias are used to connect 3 tiers.

The simulated code performance is shown in Fig. 6. The black curve shows the BER vs. SNR performance up to BER of 10^{-5} . The grey curve shows the fast iteration convergence, which yields low-power dissipation of estimated 430mw.



Figure 6. The simulated LDPC decoder performance.

Table 1 summarizes the taped out design's characteristics and a comparison 2D design. The 2D design was accomplished by putting all devices on one tier with the same technology and standard cells. We can see that the 3D implementation achieved a significant advantage over the same technology 2D implementation in terms of wire length, area, clock skew, and buffer size. The improvement in terms of power-delay-area product is more than an order of magnitude: $2.5 \cdot 3 \cdot 1.75 = 13.125$.

	2D design	3D design
area (mm*mm)	18.238*15.92=	(6.4*6.227)*3=
area (mini-min)	290.35	119.56
total wire length (m)	182 42	22.39 + 22.57 + 22.46 =
total wire keigul (iii)	102.42	67.42
max. WL before buffer	13.82	8.68
insertion (mm)	10102	0100
max. wL after buffer	4	4.17
insertion (mm)		
buffer used	32900	24636
clock skew (ns)	2.33	1
power dissipation(mw)	750	430

Table 1. The comparison between 3D and 2D designs.

V. Conclusion

A fully parallel LDPC decoder has been implemented on a 3-tier 3D IC process with 2Gb/s throughput and 430mw power consumption. The significance of this work is three-fold: (1) It is the *first* large-scale 3D ASIC implementation. (2) It is for the first time, by real silicon tape out and simulation, 3D IC process with 3-tier integration was shown to yield an order of magnitude improvement over the corresponding 2D process, in terms of power-delay-area product. (3) It is the first time that an automated 3D design flow has been developed and used to tape out a large-scale silicon ASIC design.

Acknowledgements: Dr. Guoyong Shi, Dr. Sambuddha Bhattacharya, and Dr. Lei Yang contributed to the 3D LDPC design and 3D-CAD tool development. Dr. Craig Keast, Dr. Peter Waytt, and Dr. James Burns of MIT Lincoln Lab contributed to the 3D fabrication.

References

[1] D. Mackay and R.M. Neal, "Near Shannon limit performance LDPC Codes", Electron Letters, 1996.

[2] A. Blanksby and C.J. Howland, "A 690-mw 1-Gb/s 1024b, 1/2 rate LDPC Code Decoder", JSSC, 2002.

[3] J. Burns et. al., "Three-dimensional integrated circuits for low-power, high bandwidth systems on a chip", ISSCC, 2001.



Figure 4. 3D LDPC design flow.