

Supporting Vertical Links for 3D Networks-on-Chip: Toward an Automated Design and Analysis Flow

Igor Loi, Federico Angiolini and Luca Benini, *Fellow, IEEE*
Department of Electronic Engineering and Information Science (DEIS), University of Bologna
Viale Risorgimento 2, 40136 Bologna, Italy
{iloi, fangiolini, lbenini}@deis.unibo.it

Abstract—Three-dimensional (3D) manufacturing technologies are viewed as promising solutions to the bandwidth bottlenecks in VLSI communication. At the architectural level, Networks-on-chip (NoCs) have been proposed to address the complexity of interconnecting an ever-growing number of cores, memories and peripherals. NoCs are a promising choice for implementing scalable 3D interconnect architectures. However, the development of 3D NoCs is still at an early development stage. In this paper, we present a semi-automated design flow for 3D NoCs. Starting from an accurate physical and geometric model of Through-Silicon Vias (TSVs), we extract a circuit-level model for vertical interconnections, and we use it to evaluate the design implications of extending switch architectures with ports in the vertical direction. In addition, we present a design flow allowing for post-layout simulation of NoCs with links in all three physical dimensions.

Index Terms—3D Integrated Circuits, NoCs, Wafer Bonding, Vertical Integration.

I. INTRODUCTION

OVER the years, advances in silicon technology have enabled the integration of larger and larger amounts of processing elements and memories, with increasing communication requirements at their interfaces. Simultaneously, there has been a strong push towards the mixing of functional blocks which may require some processing steps differentiating them from plain CMOS, such as DRAM, MEMS, passive and active analog circuitry, optoelectronic elements, chemical sensors, actuators, *etc.*. Vertically stacking multiple layers of silicon is an attractive way of sustaining the pace of the improvement in functionality, while providing sufficient communication bandwidth and pursuing design objectives such as small package sizes, minimum footprints and modularity.

In planar implementations, interconnects are becoming a limiting factor to achieve design closure. This is due to several issues, such as the growing ratio of wire delay *vs.* logic delay, signal integrity concerns and stringent bandwidth requirements. At the system level, the key challenge is configuring, optimizing and verifying the communication architecture across many degrees of freedom in terms of topology, architecture and interface protocols. The Network-on-Chip (NoC) paradigm, which brings packet-switching networking concepts to the on-die level, has been proposed [1], [2] to systematically tackle these challenges. NoCs are a structured, predictable and scalable approach to the problem, centered around wire segmentation and point-to-point signaling.

The simultaneous emergence of 3D integration technologies and NoCs exposes new opportunities and new challenges to system designers. The structured nature of NoCs seem to be an ideal way of encapsulating the design properties and requirements (such as heterogeneous wiring resources, large degree of parallelism, architectural heterogeneity) of three-dimensional integration. However, design tradeoffs and design technology support for 3D NoCs are yet to be explored in depth.

The first contribution of our work is the construction of a circuit-level model for vertical interconnects (Through-Silicon Vias, TSVs), based on accurate three-dimensional parasitic extraction. Comparative analysis demonstrates that not only vertical interconnects are usable, but that they are highly competitive with horizontal wires in terms of delay and power, with a reasonable area overhead. As a second main contribution, we extend a two-dimensional NoC switch architecture to deal with vertical links. Our third contribution is the development of a prototype design flow for automatic instantiation of three-dimensional NoCs. Finally, we present a case study where a planar NoC topology is folded and implemented across two chip layers.

II. PREVIOUS WORK

A. Vertical Stacking

A number of technologies for 3D chip manufacturing have been explored in recent years, including transistor stacking [3], die-on-wafer stacking [4], wafer stacking [5], chip stacking [6]. In this paper we focus on wafer stacking approaches, as one of the most promising avenues for the implementation of high-performance yet inexpensive (multiple 3D chips can be processed in a single pass) three-dimensional ICs. Wafer stacking relies on Through-Silicon Vias (TSVs) [7] for vertical connectivity, guaranteeing low parasitics (*i.e.* low latency and power) and, if needed, extremely high densities of vertical wires (*i.e.* high bandwidth). Tezzaron Semiconductor Corporation [8] and IBM Technologies [9] are active players in this field; the major differences between their processes are in wafer bonding methodologies and TSV formation. The former resorts to via formation followed by high-temperature wafer bonding, so that electrical connectivity and bonding strength are guaranteed by thermocompression. The latter uses oxide fusion bonding at room temperature, allowing a very high precision alignment, while vias are formed after the wafers have been bonded together. In this paper, we will use fabrication technology parameters disclosed in previous literature by these manufacturers [8], [9].

B. Networks-on-Chip

NoCs have been suggested as a scalable communication fabric [1], [2]. From the architectural point of view, a complete scheme is presented for example in [10], while specific topics are tackled in several works: flow control protocols [11], router power estimations [12], Quality of Service (QoS) provisions [13], [14], asynchronous implementations [15], [16], [17]. CAD tools for NoC instantiation and optimization can be found for example in [18], [19].

The synthesis flow of NoCs has been explored by several groups. Layouts are presented in [20], [21], a test chip is shown in [22], and an FPGA target is provided for [23], [24]. Synthesis and layout results for the *xpipes* library of component blocks that we will leverage upon are detailed

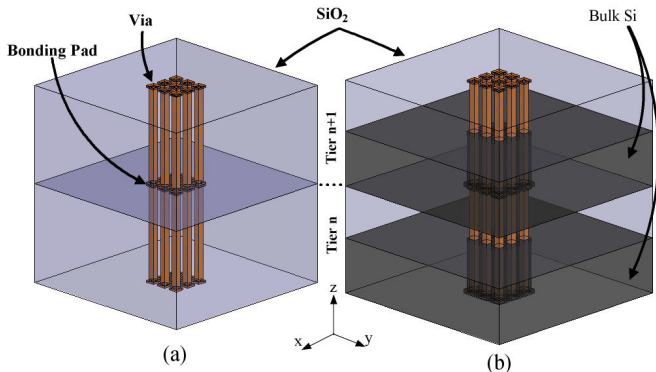


Fig. 1. Through-Silicon Vias in (a) SOI and (b) bulk-silicon technologies.

in [25], [26]. While most efforts have been aiming at standard cell ASIC targets, some groups have been doing custom design [22].

Some research is being undertaken on 3D NoCs. For example, in [27], [28] alternate ways of interconnecting 3D chips are contrasted; namely, the authors focus on several variants of 3D meshes, stacked meshes, stacked tori, *etc.*. The main focus of the authors is on topologies and on performance metrics, while the physical implementation is not studied in depth. Our work is orthogonal and complementary, as we provide accurate characterization of physical effects and parasitics, including coupling capacitances, and discuss a complete flow to implement a 3D NoC at the layout level. In [29], the authors propose a dimension decomposition scheme to optimize the cost of 3D NoC switches, and present some area and frequency figures derived from a physical implementation. The fundamental assumption of their work is that a regular, homogeneous NoC is the best solution for a 3D design, and therefore the next logical step is to reduce the cost of each required building block. However, we believe that, for such complex designs as stacked 3D chips, which are likely to mix logic layers with memory layers and even more uncommon functionality, heterogeneity will likely be significant, especially along the vertical axis. For this reason, we propose a more general approach, where the designer is allowed to choose among planar and vertical communication on a switch-by-switch basis, without any topological constraint. Post-silicon nano-scale 3D interconnections have also been recently investigated [30], but large scale availability of these technologies in the near future is uncertain. To the best of our knowledge, no previous work fully characterizes the vertical interconnections for use in NoCs, especially with respect to physical implementation and timing requirements.

III. PHYSICAL MODELING OF VERTICAL TSVs

To be useful for a NoC infrastructure, a vertical wire should not be used in isolation; instead, to simplify routing, it is better to create buses of such wires. The geometry of a TSV bus connecting adjacent stacked wafers is shown schematically in Figure 1 for two manufacturing scenarios: Silicon on Insulator (SOI) and bulk-silicon technologies. Given the physical proximity of the TSVs, concerns related to capacitive coupling within such buses may arise. In this section, we quantify the delay in a bus formed by vertical TSVs for both the SOI and bulk-silicon cases.

TSV models are obtained with the Ansoft Q3D extractor [31], a quasi-static electromagnetic-field simulation for parasitic extraction of electronic components, which utilizes finite element algorithms and the Method of Moments to compute the RLC parameters of a 3D structure. This makes the study of signal integrity (crosstalk, ground bounce) and delay possible.




Fig. 2. Schematic representation of a bundle of 3D vias.

The starting point of our analysis is a simple configuration composed of nine TSVs placed in a 3×3 grid structure. The baseline configuration we study (see Figure 2) can be summarized as:

- Copper vias
- $4\mu\text{m} \times 4\mu\text{m}$ via cross-section ($W \times L$)
- $5\mu\text{m} \times 5\mu\text{m}$ pads at via extremities
- $8\mu\text{m}$ via pitch
- $1\mu\text{m}$ oxide thickness (t_{OX}) (only for bulk silicon)
- $50\mu\text{m}$ layer thickness ($25\mu\text{m}$ bulk silicon and $25\mu\text{m}$ SiO_2)

Delay is a function of resistance and capacitance. Resistance can be described with a single parameter as a function of via length ℓ , cross-section σ and resistivity ρ :

$$R = \frac{\rho \times \ell}{\sigma} \quad (1)$$

For example, copper TSVs with $4 \times 4\mu\text{m}$ diameter show a resistance around $1.18\text{m}\Omega$ per μm . The skin effect, at these sizes, is negligible at frequencies of few GHz, and a comparison between vias and top metal wires (Metal 8, 130nm technology node) having $0.4 \times 0.8\mu\text{m}$ cross section shows that the TSV resistance per unit of length is fifty times smaller.

Capacitance, on the other hand, due to coupling effects, poses several more issues. Therefore, we resort to a capacitance matrix $\overline{\overline{C}}$ (Equation 2):

$$\overline{\overline{C}} = \begin{pmatrix} C_{1,1} & -C_{1,2} & \dots & -C_{1,n} \\ -C_{2,1} & C_{2,2} & \dots & -C_{2,n} \\ \dots & \dots & \dots & \dots \\ -C_{n,1} & -C_{n,2} & \dots & C_{n,n} \end{pmatrix} \quad (2)$$

In this matrix, the elements outside of the diagonal represent inter-via coupling, with inverted sign, while the ones along the diagonal are the sum of the capacitances towards the ground plane ($C_{i,0}$ - not explicitly reported in the matrix) plus the coupling capacitances:

$$C_{ii} = C_{i,0} + C_{i,1} + \dots + C_{i,i-1} + C_{i,i+1} + \dots + C_{i,n} \quad (3)$$

In Tables I and II we report extraction results for the capacitance of vias in SOI and bulk-silicon TSVs, respectively, for the reference case. The capacitance towards the ground plane is negligible in the SOI case, since the whole structure is "floating", but it is the dominant element in bulk-silicon technology. On the other hand, due to the presence of a passivation coating around the TSVs in the bulk-silicon case, the SOI scenario exhibits much larger coupling capacitances among the vias.

We can analyze the behavior of TSVs in different geometries using our geometric model. In Figure 3 we sweep the TSV diameter, from $0.5\mu\text{m}$ to $6\mu\text{m}$, while keeping the TSV pitch constant at $8\mu\text{m}$. Capacitance in the bulk-silicon case increases linearly with the diameter, while the increase is steeper for SOI. This is due to the fact that, in both technologies, the lateral via surface, which determines the

C [fF]	Ground	M	N	S	W	E	SW	NW	NE	SE
M	0.00	11.41	-2.43	-2.43	-2.43	-2.43	-0.41	-0.42	-0.41	-0.42
N	0.00	-2.43	10.13	-0.03	-0.47	-0.47	-0.07	-3.19	-3.18	-0.07
S	0.00	-2.43	-0.03	10.13	-0.47	-0.47	-3.19	-0.07	-0.08	-3.18
W	0.00	-2.43	-0.47	-0.47	10.13	-0.03	-3.18	-3.19	-0.07	-0.07
E	0.00	-2.43	-0.47	-0.47	-0.03	10.13	-0.08	-0.07	-3.19	-3.18
SW	0.00	-0.41	-0.07	-3.19	-3.18	-0.08	8.32	-0.40	-0.11	-0.41
NW	0.00	-0.42	-3.19	-0.07	-3.19	-0.07	-0.40	8.31	-0.40	-0.12
NE	0.00	-0.41	-3.18	-0.08	-0.07	-3.19	-0.11	-0.40	8.32	-0.41
SE	0.00	-0.42	-0.07	-3.18	-0.07	-3.18	-0.41	-0.12	-0.41	8.31

TABLE I

CAPACITANCE MATRIX OF TSVs IN SOI TECHNOLOGY. M = MIDDLE VIA; THE OTHER VIAS ARE LABELED ACCORDING TO THEIR POSITIONING WITH RESPECT TO IT (N = NORTH, etc.). "GROUND" REFERS TO THE GROUND PLANE ($C_{i,0}$).

C [fF]	Ground	M	N	S	W	E	SW	NW	NE	SE
M	-17.7	23.89	-1.20	-1.21	-1.20	-1.20	-0.33	-0.33	-0.36	-0.36
N	-18.1	-1.20	23.26	-0.09	-0.39	-0.34	-0.05	-1.58	-1.52	-0.05
S	-18.3	-1.21	-0.09	23.39	-0.35	-0.33	-1.47	-0.06	-0.05	-1.56
W	-18.1	-1.20	-0.39	-0.35	23.25	-0.09	-1.57	-1.48	-0.05	-0.05
E	-18.3	-1.20	-0.34	-0.33	-0.09	23.42	-0.05	-0.06	-1.55	-1.52
SW	-18.6	-0.33	-0.05	-1.47	-1.57	-0.05	22.23	-0.11	0.00	-0.13
NW	-18.5	-0.33	-1.58	-0.06	-1.48	-0.06	-0.11	22.16	-0.11	-0.01
NE	-18.5	-0.36	-1.52	-0.05	-0.05	-1.55	0.00	-0.11	22.24	-0.13
SE	-18.3	-0.36	-0.05	-1.56	-0.05	-1.52	-0.13	-0.01	-0.13	22.07

TABLE II

CAPACITANCE MATRIX OF TSVs IN BULK-SILICON TECHNOLOGY. M = MIDDLE VIA; THE OTHER VIAS ARE LABELED ACCORDING TO THEIR POSITIONING WITH RESPECT TO IT (N = NORTH, etc.). "GROUND" REFERS TO THE GROUND PLANE ($C_{i,0}$).

coupling, is becoming larger. Further, the distance among the lateral surfaces decreases, since the pitch is constant. However this effect is most relevant in the SOI scenario, whereas, in bulk-silicon, the passivation layer surrounding each TSV (t_{OX} thickness) dampens the increase in coupling.

It is also interesting to sweep via pitch while keeping the TSV diameter constant (e.g., at $4\mu m$). The curves are dual with respect to the previous plot, since increasing via diameters has a similar effect as decreasing via pitches. The most interesting property to be observed is the discontinuity in the bulk-silicon curves at the $6\mu m$ pitch threshold, which represents the point where two adjacent TSVs are actually in contact. This is because vias have a $4\mu m$ diameter, plus, only for the bulk-silicon case, an insulating coating $1\mu m$ thick. Below the $6\mu m$ threshold, we assume that TSVs are dug into a solid SiO_2 structure, and are therefore only separated by a thin oxide layer; above the threshold, a silicon "screen" appears in the middle as each TSV is the result of a separate etching in the silicon substrate. The presence or absence of the silicon layer changes substantially the parasitic capacitance behaviour.

The complete extracted circuit model gives maximum accuracy in electrical simulation, but good insight can be gained by modeling the delay with the well-known RC approximation:

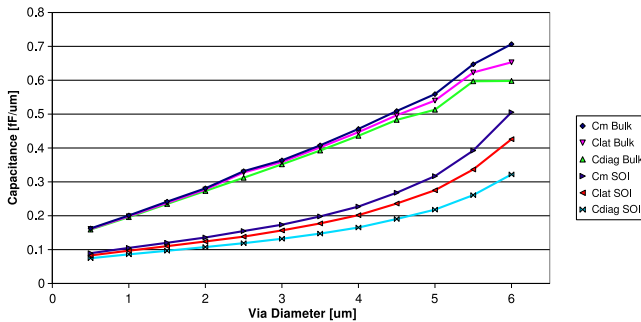


Fig. 3. Capacitance trend when sweeping the diameter of vias having a constant pitch. Figures are reported for SOI and bulk-silicon. C_m : $C_{1,1}$; C_{lat} : average of $C_{2,2}$ to $C_{5,5}$ (N, S, W, E vias); C_{diag} : average of $C_{6,6}$ to $C_{9,9}$ (SW, NW, NE, SE vias).




Fig. 4. Capacitance trend when sweeping the pitch of vias having a constant diameter. Figures are reported for SOI and bulk silicon. C_m : $C_{1,1}$; C_{lat} : average of $C_{2,2}$ to $C_{5,5}$ (N, S, W, E vias); C_{diag} : average of $C_{6,6}$ to $C_{9,9}$ (SW, NW, NE, SE vias).

$$t_D = 0.35 \times R \times C \quad (4)$$

In the formula, contact resistance and load capacitance (e.g. buffers or flip flop at the end of the line) should be taken into account. Since TSVs are interconnected by means of metal bonding, we estimate the contact resistance [32] to be $100m\Omega$ per layer. Delay estimates using Equation 4 are in good agreement with SPICE simulations. For example, $16ps$ to $18.5ps$ of delay (for SOI and bulk silicon, respectively) are found when the TSV diameter is set to $4\mu m$ and the pitch to $8\mu m$.

To put these results in perspective, the maximum un-repeated planar line length in Metal 2 and Metal 3, in the same technology, is $1.5mm$. Using a planar inter-switch link of this length as a reference, we observe that vertical links exhibit roughly one order of magnitude lower capacitive load. Roughly the same ratio can be found for resistance. As a consequence, even after taking coupling effects of tightly packed TSV bundles into account, vertical links turn out to be substantially faster and more energy efficient than moderate size planar links.

IV. INTEGRATION OF TSVs WITHIN NOC SWITCHES

NoC components and NoC design tools require modifications to support vertical links implemented with TSVs. As discussed in Section II, 3D designs are likely to expose a large degree of heterogeneity, especially along the vertical axis. Therefore, we choose to base our integration effort on the \times pipes [26] NoC library, which supports arbitrary connectivity, and on its instantiation toolchain [33]. Thus, we can leverage a semi-automatic design flow, from RTL description to layout-level verification.

\times pipes switches come in two radically different variants, conceived to best match two flow control protocols. The first is ACK/NACK, a retransmission-based protocol featuring increased error resilience. The second is STALL/GO, a simple variant of credit-based flow control allowing for pipelined links to be transparently deployed. In the ACK/NACK case, output buffers need to be inserted within switches, since any transmitted packet should be stored for potential retransmission. This implies a hardware cost, but it also means that NoC links are enclosed between two clocked buffers at the sending and receiving ends. Hence, a whole clock period is available for signal propagation along the wires of the inter-switch links. In any case, the link length and the switch logic are decoupled by the output buffer.

In contrast, in STALL/GO, low switching latency and reduced buffer cost are the main goals. \times pipes STALL/GO switches therefore adopt a lean architecture, where only switch inputs are buffered. In other words, the switch logic and the link propagation time (up to the following switch or to the first

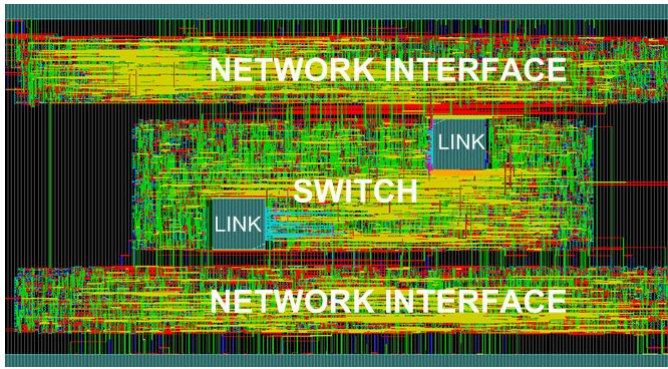


Fig. 5. Layout detail: a switch is attached to the LEF macros of two vertical links.

link pipeline stage) contribute to a the same timing path, which becomes the bottleneck for the system. While ACK/NACK transparently allows for links of arbitrary propagation time, possibly just requiring the insertion of pipeline stages, with STALL/GO the link delay directly impacts the maximum operating frequency of the switches and of the whole NoC.

We leverage the information gathered in Section III to build LEF (Library Exchange Format) descriptions of vertical vias. LEF macros are standard hardware descriptions at the layout level, including information about process technology, cell placement, routing and pins/pads. Based on these macros, TSVs can be accurately inserted within the design during the placement and routing stage; they are simply attached to the input or output pins of a switch port, just as a horizontal bus would. At the RTL level, on the other hand, the design can still be unchanged with respect to a 2D implementation. This brings several advantages: (i) the presence of vertical wires is totally transparent to the architectural and functional views of the architecture; (ii) a chip may feature any degree of connectivity heterogeneity since vertical links can be added or exchanged for horizontal ones; (iii) vertical bandwidth can be added only where needed in the chip, saving switch ports everywhere else; (iv) building upon the savings brought by the previous item, the set of switches with vertical ports, *i.e.* the ones located where vertical bandwidth is really needed, can have ideal performance because they can be implemented as full crossbars.

Thanks to this approach, a complete flow is achieved; this includes the ability to extract and simulate a 3D layout, where all switch ports are exposed to proper timing constraints and load information is available for both horizontal and vertical connections. A depiction of a sample layout featuring a 3x3 switch with vertical ports is presented in Figure 5. The arrangement of the TSV macros is the one we identified to offer the best timing requirements: close to the pinout of the switch, so as to guarantee minimum length of the wire from the switch to the base of the via, thus reducing parasitics.

The choice of a NoC topology must be performed by taking into account available performance information. Therefore, it is important to build a timing model of the switches. In Figure 6, we explore the frequency that STALL/GO and ACK/NACK switches of different cardinalities can achieve when driving horizontal (1.5mm) or vertical (50 μ m) links. As expected, ACK/NACK switches don't change operating frequency when moving to 3D structures, since their frequency bottleneck is given by the switch logic and is not affected by link performance. STALL/GO is, in general, slightly slower than ACK/NACK due to the contribution of link delay on critical paths; however, when used in combination with TSVs, it regains 30-50 MHz, *i.e.* at 50 to 75% of the frequency gap, while maintaining its low-overhead properties (and single-cycle latency). In other words, the NoC can be clocked faster when slow horizontal links are replaced by fast vertical links.

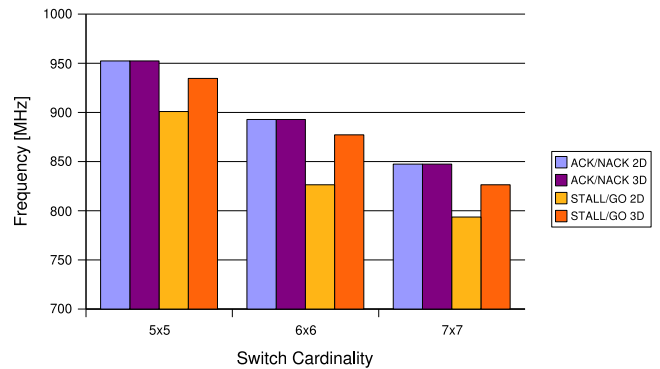


Fig. 6. Maximum frequency achievable by STALL/GO vs. ACK/NACK switches in 2D and 3D flows, for varying switch cardinalities.

V. IMPLEMENTATION OF TSV-BASED NOCS

As a validation of our flow, we present a NoC implementation based on a 2D 3x2 quasi-mesh (called simply mesh in the following) and migrate it to a 3D arrangement (Figures 7 and 8). The 3D mapping is achieved by splitting in two halves the mesh and overlapping them in separate chip layers, with communication achieved through TSVs. The stacked topology has exactly the same functionality of the two-dimensional implementation.

As a first step, we leverage SunFloor [33] to instantiate the 2D mesh. There is no need to modify the RTL output of SunFloor in any way. Next, we identify the best partitioning for mapping onto the layer stack. This task is, at present, done manually, due to the large set of constraints involved. These include manufacturing limitations, chip pinout, area considerations, bandwidth demands, thermal requirements, *etc.*. For example, our test 3x2 mesh connects three processors and three memories; since we assume that processors cannot be stacked on top of each other, to avoid the formation of hot spots, we interleave processors and memories. \times pipes links connect either two different switches or a switch and a network interface; our choice is to cut two-dimensional topologies across switch-to-switch links, replacing the latter with an upstream and a downstream port.

Then we perform synthesis, placement and routing of the RTL in two separate runs, one per design partition. During placement, we insert TSV macros at the proper switch boundaries. We choose minimum TSV diameter (4 μ m) and pitch achievable in current technologies. The area overhead of each TSV is 64 μ m² (8x8). For each bidirectional vertical switch port (*e.g.* the Up one) we have $2 \times (5 + DataWidth)$ TSVs, where the factor 2 is due to the presence of one input and one output port for bidirectionality, 5 is the number of control signals, and *DataWidth* is the width (in bits) of the inter-switch data link. In the example of a 6x6 switch and assuming a *DataWidth* of 28 bits, the area overhead is about 6% for ACK/NACK and 9% for STALL/GO. In exchange for this small area cost, switches can operate around 10% faster and less buffering can be deployed (saving up to 13% of the sequential area).

VI. CONCLUSIONS AND FUTURE WORK

In this work, we have studied the performance and system-level impact of through-silicon vias as one of the possible ways to implement high-density vertical NoC links. We have shown that, even when accounting for the coupling effects in dense vertical link bundles, the parasitics associated with TSVs are one order of magnitude smaller than traditional horizontal wires, making 3D NoCs a very promising approach. We have shown how to design NoC switches with vertical ports. Finally, we have shown that our semi-automated flow is capable of




Fig. 7. 2D 3x2 mesh NoC topology and one possible 3D re-implementation.

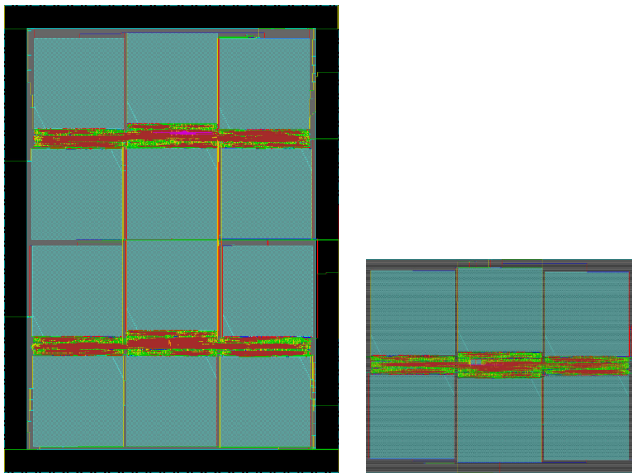


Fig. 8. Layouts for (a) the 2D 3x2 mesh, (b) one of the halves of its 3D re-implementation.

generating layouts of 3D NoCs which are fully compatible with accurate post-layout timing, area and power analysis.

Research on 3D NoCs is just now beginning, and much work remains to be done. Among the areas requiring more attention, we plan on focusing on design partitioning for 3D mapping, and on the issue of how to build 3D NoCs which can efficiently connect layers running at different clock frequencies.

ACKNOWLEDGMENTS

This work is supported by a grant from Semiconductor Research Corporation (SRC project number 1188) and a grant by STMicroelectronics for DEIS.

REFERENCES

- [1] W. J. Dally and B. Towles, "Route packets, not wires: On-chip interconnection networks," in *Proceedings of the 38th Design Automation Conference*, June 2001, pp. 684–689.
- [2] L. Benini and G. De Micheli, "Networks on chips: A new SoC paradigm," *IEEE Computer*, vol. 35, no. 1, pp. 70–78, January 2002.
- [3] B. Rajendran, R. S. Shenoy, D. J. Witte, N. S. Chokshi, R. L. DeLeon, and G. S. Tompa, "Cmos transistor processing compatible with monolithic 3-d integration," in *Proc. VLSI Interconnection (VMIC)*, 2005, pp. 76–82.
- [4] Ziptronix, *Ziptronix target vertical scalability*, 2005.
- [5] S. Christiansen, R. Singh, and U. Gosele, "Wafer direct bonding: From advanced substrate engineering to future applications in micro/nanoelectronics," in *Proceedings of the IEEE*, December 2006, pp. 2060–2106.
- [6] K. Lee, "Wafer-stacked package technology for high-performance system," in *RTI Int. technology Venture Forum*, 2005.
- [7] S. Spiesshoefer and et al., "Z-axis interconnects using fine pitch, nanoscale through-silicon vias: Process development," in *Electronic Components and Technology Conference*, 2004.
- [8] R. S. Patti, "Three-dimensional integrated circuits and the future of system-on-chip designs," *Proceedings of the IEEE*, vol. 94, no. 6, June 2006.

- [9] A. W. Topol, J. D. C. La Tulipe, L. Shi, D. J. Frank, K. Bernstein, S. E. Steen, A. Kumar, G. U. Singco, A. M. Young, K. W. Guarini, and M. Jeong, "Three-dimensional integrated circuits," *IBM Journal of Research and Development*, vol. 50, no. 4/5, pp. 491–506, July/September 2006.
- [10] F. Karim, A. Nguyen, S. Dey, and R. Rao, "On-chip communication architecture for OC-768 network processors," in *Proceedings of the Design Automation Conference (DAC)*, 2001, pp. 678–683.
- [11] A. Pullini, F. Angiolini, D. Bertozzi, and L. Benini, "Fault tolerance overhead in network-on-chip flow control schemes," in *Proceedings of the 18th Annual Symposium on Integrated Circuits and System Design (SBCCI)*, 2005, pp. 224–229.
- [12] W. Hang-Sheng, Z. Xiping, P. Li-Shiuan, and S. Malik, "Orion: a power-performance simulator for interconnection networks," in *Proceedings of 35th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE/ACM, November 2002, pp. 294–305.
- [13] E. Bolotin, I. Cidon, R. Ginosar, and A. Kolodny, "QNoC: QoS architecture and design process for network on chip," in *Journal of Systems Architecture*. Elsevier, 2004.
- [14] D. Wiklund and D. Liu, "SoCBUS: Switched network on chip for hard real time embedded systems," in *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS03)*. IEEE, 2003.
- [15] T. Bjerregaard and J. Sparsø, "Scheduling discipline for latency and bandwidth guarantees in asynchronous network-on-chip," in *Proceedings of the 11th IEEE International Symposium on Asynchronous Circuits and Systems (ASYNC)*, 2005, pp. 34–43.
- [16] A. Sheibanyrad, I. M. Panades, and A. Greiner, "Systematic comparison between the asynchronous and the multi-synchronous implementations of a network on chip architecture," in *Design, Automation & Test in Europe Conference & Exhibition*, April 2007, pp. 1–6.
- [17] S. Furber and J. Bainbridge, "Future trends in soc interconnect," in *Proceedings of the International Symposium on System-on-Chip (SoC)*. IEEE Computer Society, 2005.
- [18] S. Murali, M. Coenen, A. Radulescu, K. Goossens, and G. D. Micheli, "Mapping and configuration methods for multi-use-case networks on chips," in *Proceedings of the 2006 conference on Asia South Pacific design automation (ASP-DAC)*. New York, NY, USA: ACM Press, 2006, pp. 146–151.
- [19] K. Srinivasan and K. Chatha, "A methodology for layout aware design and optimization of custom network-on-chip architectures," in *Proceedings of the 7th International Symposium on Quality Electronic Design (ISQED)*. IEEE Computer Society, 2006.
- [20] A. Radulescu, J. Dielissen, K. Goossens, E. Rijpkema, and P. Wielage, "An efficient on-chip network interface offering guaranteed services, shared-memory abstraction, and flexible network configuration," in *Proceedings of the 2004 Design, Automation and Test in Europe Conference (DATE)*. IEEE, 2004.
- [21] A. Andriahantenaina and A. Greiner, "Micro-network for SoC: Implementation of a 32-port SPIN network," in *The Proceedings of Design, Automation and Test in Europe Conference and Exhibition*. IEEE, 2003, pp. 1128–1129.
- [22] K. Lee, S.-J. Lee, S.-E. Kim, H.-M. Choi, D. Kim, S. Kim, M.-W. Lee, and H.-J. Yoo, "A 51mW 1.6GHz on-chip network for low-power heterogeneous SoC platform," in *Digest of Technical Papers of the 2004 IEEE International Solid-State Circuits Conference (ISSC)*. IEEE Computer Society, 2004, pp. 152–158.
- [23] C. A. Zeferino and A. A. Susin, "SoCIN: A parametric and scalable network-on-chip," in *Proceedings of the 16th Symposium on Integrated Circuits and Systems Design (SBCCI03)*, 2003, pp. 34–43.
- [24] P. T. Wolkotte, P. K. Holzspies, and G. J. Smit, "Fast, accurate and detailed noc simulations," in *Proceedings of the First International Symposium on Networks-on-Chip (NOCs)*. IEEE Computer Society, 2007.
- [25] F. Angiolini, P. Meloni, D. Bertozzi, L. Benini, S. Carta, and L. Raffo, "Networks on chips: A synthesis perspective," in *Proceedings of the 2005 ParCo Conference*, 2005.
- [26] F. Angiolini, P. Meloni, S. Carta, L. Raffo, and L. Benini, "A layout-aware analysis of networks-on-chip and traditional interconnects for mpsoCs," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 26, no. 3, pp. 421–434, March 2007.
- [27] V. F. Pavlidis and E. G. Friedman, "3-D topologies for networks-on-chip," in *Proceedings of the IEEE SoC Conference (SOCC)*. IEEE Computer Society, 2006, pp. 285–288.
- [28] B. Feero and P. P. Pande, "Performance evaluation for three-dimensional networks-on-chip," in *Proceedings of the IEEE Annual Symposium on VLSI (ISVLSI)*. IEEE Computer Society, 2007, pp. 305–310.
- [29] J. Kim, C. Nicopoulos, D. Park, R. Das, Y. Xie, N. Vijaykrishnan, M. S. Yousef, and C. R. Das, "A novel dimensionally-decomposed router for on-chip communication in 3d architectures," in *Proceedings of the 34th International Symposium on Computer Architecture (ISCA)*, 2007.
- [30] S. Fujita, K. Nomura, K. Abe, and T. Lee, "3d on-chip networking technology based on post-silicon devices for future networks-on-chip," in *Nano-Networks and Workshops*, September 2006, pp. 1–5.
- [31] A. Corp., "Q3d extractor," 2007, http://www.ansoft.com/products/si/q3d_extractor/.
- [32] K. N. Chen, A. Fan, and C. S. T. ans R. Reif, "Contact resistance measurement of bonded copper interconnects for three-dimensional integration technology," *IEEE ELECTRONIC DEVICE LETTERS*, vol. 25, no. 1, January 2005.
- [33] S. Murali, P. Meloni, F. Angiolini, D. Aienza, S. Carta, L. Benini, and G. D. Micheli, "Designing application-specific networks on chips with floorplan information," in *Proceedings of the 2006 International Conference on Computer-Aided Design (ICCAD)*. New York, NY, USA: ACM Press, 2006, pp. 355–362.