

A Low-overhead Fault Tolerance Scheme for TSV-based 3D Network on Chip Links

Igor Loi[†], Subhasish Mitra[‡], Thomas H. Lee[‡], Shinobu Fujita^{*} and Luca Benini[†]

[†]DEIS, University of Bologna, Bologna, Italy

[‡]Stanford University, California, Usa

^{*}Toshiba, San Jose, CA, USA (Kawasaki, Kanagawa, Japan)

igor.loi@unibo.it, subh@stanford.edu, tomlee@smirc.stanford.edu, shinobu.fujita@toshiba.co.jp
and lbenini@deis.unibo.it

Abstract— Three-dimensional die stacking integration provides the ability to stack multiple layers of processed silicon with a large number of vertical interconnects. Through Silicon Vias (TSVs) provide a promising area- and power-efficient way to support communication between different stack layers. Unfortunately, low TSV yield significantly impacts design of three-dimensional die stacks with a large number of TSVs. This paper presents a defect-tolerance technique for TSVs-based multi-bit links through an efficient and effective use of redundancy. This technique is ideally suited for three-dimensional network-on-chip (NoC) links. Simulation results demonstrate significant yield improvement, from 66% to 98%, with a low area cost (17% on a vertical link in a NoC switch, which leads a modest 2.1% increase the total switch area) in 130nm technology, with minimal impact of VLSI design and test flows.

I. INTRODUCTION

FOR future integrated system design, two major trends are emerging. Communication-centric architectures based on the Network on Chip (NoC) design paradigm [1], [2] to tackle interconnect and architectural scalability challenges. Three-Dimensional Integrated Circuits (3DICs) that provide a promising technological solution to alleviate the interconnect, I/O bandwidth and latency bottlenecks.

3DICs may enable heterogeneous integration and new classes of applications through significantly improved performance and energy efficiency of complex system architectures (e.g. technologies from Tezzaron Semiconductor Corporation [3], IMEC, MIT Lincoln Labs, and IBM [4]). One of the most promising technologies for 3D integration is based on Through Silicon Vias (TSVs), which cut across thinned silicon substrates to establish inter-die connectivity after die-bonding.

Three-Dimensional Network on Chips (3DNoCs) combine the benefits of short vertical interconnects of 3DICs and the scalability of NoCs. 3DNoCs support both horizontal and vertical links. A vertical link can be physically implemented as a cluster of TSVs. TSVs allow fine pitch, high density and high compatibility with the standard CMOS process. Unfortunately, currently available processes for TSV fabrication have relatively low yield (compared to standard 2D processes). Figure 1 shows limited yield of TSVs from three different process technologies: HRI [5], IMEC [6] and IBM [7].

In this paper, we describe the design of a defect-tolerant TSV-based multi-bit vertical link which enables significant yield improvement with respect to random (complete or partial) open defects at an extremely low cost. Like traditional

defect-tolerance techniques (such as those used for memories), our technique also relies on redundancy. Our major contribution is in a simple and efficient design of such a defect-tolerant TSV-based link at lowest cost, and also with minimal impact on the overall integrated system design and production test flows. While this TSV-based link design is generally applicable for both NoC-based and bus-based 3D interconnects, it is especially useful for 3DNoCs because it takes advantage of the NoC switch architecture to introduce minimal system-level area impact.

The main contributions of this paper are:

- Introduction of a robust, defect-tolerant, vertical link architecture (for TSVs) to overcome challenges of low yield for current TSV fabrication processes;
- Integration of the defect-tolerant 3D link into a complete three-dimensional Network on Chip design flow;
- Experimental evaluation, performed at the layout level, including full placement and routing, to evaluate benefits, feasibility and hardware costs.

In our experiments, we achieve significant yield improvements (from 66% to 98% for 4.2M TSVs design, arranged in 100K spots made up of 42 TSVs each) for random (complete and partial) open defects that pose major challenges for TSVs. Our layout results demonstrate the feasibility of this approach and its low cost (17% on a vertical link in a NoC switch, which leads a modest 2.1% increase in the switch area).

II. RELATED WORK

Interconnect scaling has become one of the most crucial challenges in chip design, and is expected to get worse in the




Fig. 1. Yield trend for TSVs in three different processes: IBM, HRI and IMEC. Only random (complete or partial) open defects are considered in this figure, since misalignments are well controlled during the bonding phase. Yield is evaluated using the Poisson distribution.




Fig. 2. Cross-section of a vertical link across two tiers. The figure also shows the worst-case misalignment scenario

future. 3D integration and Network on Chip design methodologies are expected to overcome many of these challenges. NoCs have been suggested as a scalable communication fabric [1], [2]. 3D integration has been proposed in different ways (e.g. Tezzaron Semiconductor Corporation [3], IMEC, MIT Lincoln Labs, and IBM Technologies [4]) providing promising solutions to enable connectivity along the vertical direction.

Recently, some research has been undertaken on 3DNoCs. In [8], 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. Post-silicon nano-scale 3D interconnections have also been recently investigated [9], [10], but large scale availability of these technologies in the near future is uncertain.

As technology scales, fault tolerance is becoming a key concern in on-chip communication. Optical Proximity Correction (OPC) and redundant via placement [11] have solved a huge number of cases of faults related, mainly, to interconnects.

Recent experiments by HRI on 3DICs report very high yields of over 60%, and the redundancy scheme used realizes each vertical interconnect as a pair of vias (twins) [5].

Despite the research undertaken on 3DICs and recently on 3DNoCs, to date, yield improvements for vertical links of 3DNoCs have never been studied. In this paper, we propose a novel scheme to overcome this limitation. The starting point of this work is [12], [13], where a thorough physical and timing analysis of the vertical links has been conducted on a real 3DNoC. Further, it is worth stressing that the proposed scheme can also be applied successfully to alternative interconnection schemes, such as buses.

III. PHYSICAL LEVEL MODELING AND ANALYSIS OF TSV FAULT IMPACT

In this paper we focus on the wafer stacking approach since it is very promising for the implementation of high-performance yet inexpensive 3DICs. Wafer stacking relies on Through-Silicon Vias (TSVs) [14] for vertical connectivity, guaranteeing low parasitics (i.e. low power and propagation delay) and, if needed, extremely high densities of vertical wires (i.e. high bandwidth-per-area ratio). The electrical connectivity between different tiers is provided by creating bonding pads on the wafer surface, and then performing bonding by mechanical thermo-compression.




Fig. 3. TSVs and global wire electrical model for two stacked vias (refer to Figure 2)

The primary failure mechanisms for TSVs are misalignment and random (complete or partial) open defects [15]. Misalignment refers to unsuccessful wafer alignment prior to and during wafer bonding process (Figure 2), and is caused by shifts of bonding pads with respect to their nominal positions. Random defects comprise a variety of unpredictable physical phenomena related to the thermal compression process used in wafer stacking.

Starting from these considerations and based on [12], we have conducted a detailed study to quantify the impact of TSV failures on overall chip yield. We use an electrical model of TSVs and the bonding mechanisms for this purpose. Figure 3 shows the electrical model of two stacked vias (rendered as a T network). The vias are driven by one inverter followed by a stretch of planar interconnect (global routing). The contact resistance is related to the quality and area of bonding.

In case of misalignments (e.g. top wafer shifts along the X or Y axes or a small rotation), the bonded area decreases. This phenomenon has been modeled as a variable resistance (central resistor in Figure 3) between the two T networks, and the outcome is summarized in Table I. As can be seen, misalignments of even noticeable entity do not normally compromise functionality (which is dominated by the overall planar routing parasitics [12]) and have a minimum impact on delay. Extreme misalignment, like in the last row of table I are highly unlikely in state-of-the-art wafer bonding processes [3], [4], [16]. This motivates special emphasis on workarounds for the other main source of yield losses: random defects.

Random (complete or partial) open defects affect single vias or a small area of the interface because of failure mechanisms such as dislocations, O_2 trapped on the surface, void formation, or even mechanical failures in TSVs [17], [18], [18], [4], [19]. To model the effects of these defects, we assumed a uniform TSV defect distribution and performed several Monte Carlo simulations. Based on our results (Section V), we concluded that random (complete or partial) open defects are far more relevant compared to misalignment problems. For this reason, we focus on these defects in the following sections.

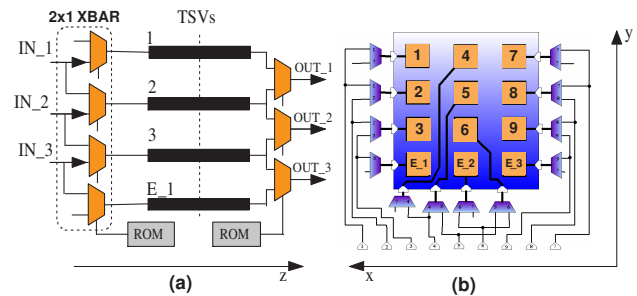


Fig. 4. Redundant Routing scheme. (a) shows a simplified crossbar scheme for dynamic routing (functional scheme). (b) shows the TSVs obstruction and the routing crossbar (the orange squares are the TSV pads). Extra pads (E.1 E.2 ...) are spread around the TSV cluster, simplifying fault bypassing by means of a 2X multiplexers.

Misalignment [μm] in X-Y	Contact Area [μm^2]	Contact Resistance [Ω]	Δ Delay [%]
0	4x4	10m	0
1	3x3	19m	< 1%
2	2x2	40m	< 1%
3	1x1	160m	< 1%
3.98	0.02x0.02	1K	22%

TABLE I

PAD CONTACT RESISTANCE AND DELAY INCREASING FOR CU-CU WAFER METAL BONDING UNDER DIFFERENT MISALIGNMENT CASES [17], [20]

IV. YIELD ENHANCEMENTS FOR 3DNoCs

In this section, we describe the target 3DNoC [12], [13] used for our experiments, and present our defect-tolerant solution for TSV-based vertical link design. As pointed out earlier, our solution can be applied not only to 3DNoCs, but also, more generally, to regular structures such as buses.

A. The reference NoC architecture

To make our study realistic, we developed our approach within the xpipes [12], [21], [13] NoC library. To enable our NoC for 3D technology we extended the xpipes switches by adding a couple of vertical ports, and we developed hard macros for the TSVs obstruction [12]. Vertical links are unidirectional, and are composed (as planar links) of data and flow control signals, traveling in opposite directions. For this work we selected a data width of 32 bit, therefore for a pair of 3D links, 76 different signals are needed overall.

B. Yield Enhancement Approaches

Among the numerous techniques to increase wafer yield of VLSI designs, we focus on hardware redundancy, deployed at design time, with some amount of post-manufacturing configuration. We use active redundancy in the form of spare pads and reconfigurable routing hardware in order to minimize the overall complexity, while gaining maximum benefits in term of efficiency (Figure 4).

The dynamic routing solution is designed to leverage post-manufacturing configurability of the TSV interconnect map. This allows us to achieve high yield while minimizing the overhead in terms of the number of pads and extra logic. Combining testing resources (e.g., scan chains¹) with such reconfigurability plays the key role in achieving yield. This solution allows us to test each vertical interconnect and diagnose defects, to isolate any failed TSV, and finally to restore functionality through reconfiguration by routing the affected signals over to the spare pads.

As we see in Figure 4 (a), in our proposed Dynamic Routing scheme, all pads are driven by a 2x1 crossbar, and each signal can be routed to two different TSVs. We explore configurations with one extra pad for each cluster (i.e. for each pad column). The crossbar is extremely small, as a strategic choice to keep the area overhead as low as possible - for each additional re-routing degree of freedom, the crossbar radix increases by a factor of one. With this lean architecture, faults are recovered by shifting affected signals to the neighboring pads, and further shifting the displaced connections over to other adjacent pads until all connections are across safe electrical structures. To clarify the recovery scheme, we shall consider Figure 4 (b).

¹The use of scan chains does not normally imply any extra cost, as they are typically integrated in every design

Supposing that pad 2 is affected by some defects (resulting e.g. in an open circuit), we route signal 3 normally through its associated pad 3, while signal 2 gets rerouted through pad 1, and therefore signal 1 gets remapped to pad E.1. Signals outside this column are not shifted since the defect is contained inside the first cluster; the recovery process is performed locally. The proper routing information is elaborated off-chip (to minimize hardware complexity and overhead) during chip testing, and is then stored on-chip into a small One Time Programmable (OTP) memory (e.g. a fuse ROM).

The importance of the testing stage is evident, as it determines all the necessary inputs to correctly set the crossbar up. To test the physical interconnect, we reuse the scan chains which are normally inserted anyway in the design, thus incurring no overhead for this. Figure 5 illustrates the hardware facilities used to test the TSVs. The TSVs are tested by injecting Test Vectors (TVs) in one tier (e.g. the bottom one). The TV is propagated to the destination tier (e.g. the top one), where it is captured and transmitted off-chip. In summary, the approach is split into five steps:

- 1) Inject test vectors (e.g. bottom tier);
- 2) Propagate test vectors across TSVs and capture them (e.g. top tier);
- 3) Scan out the captured data (e.g. top tier);
- 4) Elaborate off-chip the interconnect map;
- 5) Reconfigure the crossbar (both bottom and top tier);

The process can be performed at any speed allowed by the external I/O pins. Since the interconnect map is devised off-chip, minimal logic is required on-chip for the mapping procedure - mostly, the OTP memory to store the crossbar configurations.

V. EXPERIMENTAL RESULTS

A. Yield and Hardware Cost of the Redundant Solutions

The alternative solutions, and a non-redundant baseline case, have been synthesized with the UMC 130nm technology library and inserted into the floorplan for a 3D chip stack. Placement, routing and post layout verification have been performed. As depicted in Figure 8, the planar topology has been partitioned in two parts (dotted line), between the central routers. The topology under test (see Figure 8) includes six processors and six memories, placed on two layers. Vertical communication is achieved through the two central switches

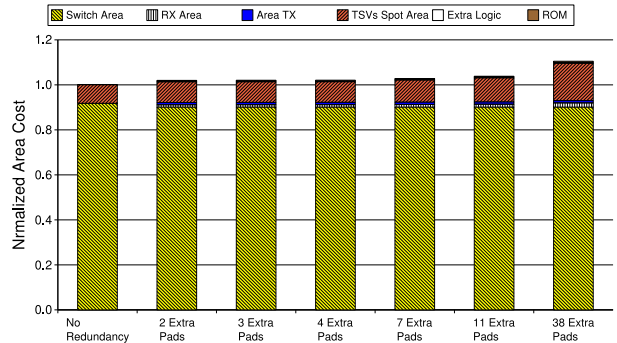


Fig. 6. Normalized area cost in case of No Redundancy and Dynamic Routing with 2, 3, 4, 7, 11 and 38 extra pads. The main contribution of this paper is resumed starting from the 2nd bar, which shows only 1.6% area overhead for 2 extra pads, 2.1% for 4 extra pads and 10.5% for full redundancy (38 extra pads)




Fig. 5. TSV NoC Test Environment: in test mode, test vectors are injected from the Test Access Point (1*) into the switch input buffer (scan), then the path through the crossbar is enabled (1*) and flow control is disabled. After some cycles the stimuli reach the next tier where they are captured (2*) from the input buffer, and then shifted out through the TAP (3*). This stream is analyzed off-chip then, based upon the failure map the OTP memories are programmed (5*), reconfiguring the crossbar to isolate failed structures

which act as a gateway for 3DNoC traffic. The reconfigurable crossbars have been inserted between the TSV pads and the switch. For a 32-bit link, the NoC protocol uses 38 bits, where the remaining 6 bits belong to flow control signaling and mesochronous handling (i.e. the clock and reset signals which are forwarded along with the data).

The nature of the reference NoC switches, namely their flow control, have influenced the adopted testing solution. During testing, a portion of the hardware works in scan mode (inject) and the other in capture mode; the flow control has to be explicitly managed to avoid the formation of communication stalls. Four scan chain groups have been inserted, driven by a simple Finite State Machine (FSM), accomplishing high efficiency and reliability. The overhead of this approach is mainly due to the crossbar logic around the via bundles, to the OTP memory and to the small FSM. The scan chain cost is not taken into account since, as mentioned before, the design must be testable anyway, and this contribution is present as well on planar ICs.

Several experiments have been conducted, especially with the dynamic routing technique, in order to evaluate how many extra pads and area may be needed for implementation, and in order to explore the trade-offs between yield and cost. We implemented six different configurations, respectively with 2, 3, 4, 7, 11 and 38 extra pads. It is worth noting that, in each unidirectional link of 38 signals, spare pads are separately needed for incoming (mostly, flow control) and outgoing (mostly, data) wires; hence the need for at least 2




Fig. 7. Yield improvement over seven different hardware configurations: no-redundancy, 2, 3, 4, 7, 11 and 38 extra pads, which correspond to 38, 40, 41, 42, 45, 49, and 76 TSVs per 3D link. A fixed defect frequency of 9.75 Defects Per Million Opportunities (DPMO) is assumed, and 4.2M TSVs design has been analyzed.




Fig. 8. 3D NoC topology. Dash boxes indicate the resources involved in the TSV test process.

spares. The latter group typically features many more wires than the former (35 vs. 3 in our example), so the correction performance is maximized with an asymmetric assignment of spares to the two groups. For example, with only 2 extra pads, no choice is available; there is only one spare for 35 outgoing signals, while the 3 incoming wires share the second spare. With 4 spares, the optimal arrangement is to assign 3 to the outgoing bundle, and the fourth to the incoming bundle. In the extreme case of 38 spares, each TSVs has a backup.

Figure 7 illustrates the yield improvement in case of 2, 3, 4, 7, 11 and 38 extra pads and based on experimental data, assuming a fixed defect frequency of 9.75 Defect Per Million Opportunities (HRI TSV process [5]). We emulated 100K TSVs links with and without redundancy. Without post-manufacturing processing, the system is unable to recover damaged vias, and tolerates only small misalignments, thus exhibiting a yield of only 68%. When Dynamic Routing redundancy is adopted, the recovery algorithm shows excellent results, especially with 2 to 7 extra pads. Further increasing the number of extra pads brings minimal yield benefits, and the increase in cost of TSV obstructions, TSV crossbar and the OTP memory may be unjustified. With only four extra pads per 3D link, yield increases from 68% to 98%.

Concerning the silicon cost, Figure 6 shows the normalized area cost in case of different degrees of redundancy applied to a single 3D link. As the number of extra pads increases, the TSVs spot and the routing logic grow in a linear fashion. The increasing area, with reference to a baseline composed of the Switch and the non-redundant TSVs link is 1.6% in case of 2 extra pads, 2.1% on 4 extra pads, and 10.5% in case of 38

extra pads. As a stand-alone component, the redundant links with 2 extra pads impacts for the 17% on a non-redundant Link. The physical implementation of the redundant hardware is depicted in Figure 9, where a pair of 3D links (42 TSVs each) is surrounded by the routing crossbar, guaranteeing low latency and better area utilization.

To evaluate the impact of the Dynamic Routing solution uses advanced technology nodes, we performed an experiment using 65nm technology library. As we can see in Table II, by scaling the technology the Dynamic Routing logic scales as well. But, the TSV obstructions show the same area, since we conservatively assumed that the TSV process is independent from the technology node used for the 2D chip and it does't scale. Therefore, the area overhead of our solution increases from 2.1% on 130nm to 3.8% on 65nm, which is still very affordable.

VI. CONCLUSIONS AND FUTURE WORK

3DICs, especially those based on Through-Silicon Vias, are gaining traction as a workaround against the increasing costs of chip miniaturization. However, the manufacturing technology is not mature enough, resulting in issues such as misalignments and random defects. Misalignment-reduction techniques have undergone significant improvements, so that today random defects must be considered the main source of yield losses. For this reason, minimizing their impact is crucial. In this paper, we study some baseline redundancy schemes and we notably propose a novel Dynamic Routing approach. The latter scheme is based on post-manufacturing study and reconfiguration of the electrical resources, leveraging a small amount of on-chip spares. The scheme proves capable of yields up to 98% with a minimum silicon cost of just 17% per TSV link in 130nm. This cost is further projected to decrease to just 12% in the newest 65nm technologies.

Future work may revolve around timing faults, which are an often underestimated source of failures.

ACKNOWLEDGMENTS

This work is the result of a strict collaboration between University of Bologna, Toshiba and Stanford Center for Integrated Systems. This work is supported by European project 214364 ICT-GALAXY 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 chip: a new SoC paradigm," *IEEE Computer*, vol. 35, no. 1, pp. 70–78, January 2002.
- [3] 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.

	130nm				
	SW Area	TSVs Area	Routing HW	Total Area	Total Area increase
NoRedundancy	54000	4864	-	58864	-
With Redundancy	53000	5376	1713	60090	+1225 (2.1%)
	65nm				
	SW Area	TSVs Area	Routing HW	Total Area	Total Area increase
NoRedundancy	13500	4864	-	18364	-
With Redundancy	13250	5376	430	19056	+692 (3.8%)

TABLE II

SILICON COST [μm^2] OF TSV REDUNDANCY SOLUTION WITH 4 EXTRA PADS IN 2 TECHNOLOGY NODES




Fig. 9. Layout detail of the bottom tier (3DICs) with emphasis on TSVs guide and configurable crossbar

- [4] 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.
- [5] N. Miyakawa, T. Maebashi, N. Nakamura, S. Nakayama, E. Hashimoto, and S. Toyoda, *New Multi-Layer Stacking Technology and Trial Manufacture*, November 2007, honda Research Institute Japan Co. Ltd.
- [6] B. Swinnen, W. Ruythooren, P. D. M. L. Bogaerts, L. Carbonell, K. D. Munck, B. Eyckens, S. Stoukatch, Tezcan, D. Sabuncuoglu, Z. Tokei, J. Vaes, J. V. Aelst, and E. Beyne, "3d integration by cu-cu thermo-compression bonding of extremely thinned bulk-si die containing 10 um pitch through-si vias," pp. 1–4, 2006.
- [7] A. W. T. et al, "Enabling soi based assembly technology for three dimensional integrated circuits," pp. 352–355, IEDM 2005.
- [8] J. Kim, C. Nicopoulos, D. Park, R. Das, Y. Xie, N. Vijaykrishnan, M. S. Yousif, 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.
- [9] 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.
- [10] S. Fujita, K. Nomura, K. Abe, and T. H. Lee, "3-d nanoarchitectures with carbon nanotube mechanical switches for future on-chip network beyond cmos architecture," *IEEE Transactions on Circuits and Systems Part I: Regular Papers*, vol. 54, no. 11, pp. 2472–2479, November 2007.
- [11] M. Rencher and F. Schellenberg, "Why interconnect and lithography modeling impacts yield," in *What's Yield got to do with IC*.
- [12] I. Loi, F. Angiolini, and L. Benini, "Supporting vertical links for 3d networks-on-chip: Toward an automated design and analysis flow," in *Proceedings of the Nano-Net Conference 2007*, 2007, pp. 23–27.
- [13] I.Loi, F.Angiolini, and L.Benini, "Developing mesochronous synchronizer to enable 3d nocs," in *Proceedings of the Date Conference 2008*, 2008, pp. 1414–1419.
- [14] S.Spiesshoefer and et al, "Z-axis interconnects using fine pitch, nanoscale through-silicon vias: Process development," in *Electronic Components and Technology Conference*, 2004.
- [15] R. Patti, "Impact of wafer-level 3d stacking on the yield of ics," in *Future Fab Intl*, September 2007. [Online]. Available: http://www.future-fab.com/documents.asp?d_id=4415
- [16] N.Miura, D.Mizoguchi, M.Inoue, T.Sakurai, and T.Kuroda, "A 195-gb/s 1.2-w inductive inter-chip wireless superconnect with transmit power control scheme for 3-d-stacked system in a package," *IEEE journal of solid state circuits*, vol. 41, no. 1, pp. 23–34, january 2006.
- [17] K.-N. Chen, A. Fan, and R. Reif, "Microstructure examination of copper wafer bonding," in <http://www-mtl.mit.edu/reif/papers/2001-knchen-JEM-manuscript.pdf>.
- [18] K.-N. Chen, C. Tan, A. Fan, and R. Reif, "Morphology and bond strength of copper wafer bonding," in *Electrochem. Solid-State Lett.*, vol. 7, 2004, pp. 14–16.
- [19] A. Papanikolaou, M. Miranda, H. Wang, F. Cathoor, M. Satyakiran, P.Marchal, B. Kaczer, C. Bruynseraede, and Z. Tokei, "Reliability issues in deep deep sub-micron technologies: time-dependent variability and its impact on embedded system design," pp. 342–347, IEDM 2006.
- [20] K. N. Chen, A. Fan, and R. Reif, "Interfacial morphologies and possible mechanisms of copper wafer bonding," in <http://www-mtl.mit.edu/users/reif/papers/2002-knchen-JMS-manuscript.pdf>.
- [21] 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.