

3D Integration for Power-Efficient Computing

D. Dutoit, E. Guthmuller, I. Miro-Panades
CEA-LETI, MINATEC Campus
38000 Grenoble, France
denis.dutoit@cea.fr

Abstract— 3D stacking is currently seen as a breakthrough technology for improving bandwidth and energy efficiency in multi-core architectures. The expectation is to solve major issues such as external memory pressure and latency while maintaining reasonable power consumption. In this paper, we show some advances in this field of research, starting with memory interface solutions as WIDEIO experience on a real chip for solving DRAM accesses issue. We explain the integration of a 512-bit memory interface in a Network-on-Chip multi-core framework and we show the performance we can achieve, these results being based on a 65nm prototype integrating 10 μ m diameter Through Silicon Vias. We then present the potentiality of new fine grain 3D stacking technology for power-efficient memory hierarchy. We expose an innovative 3D stacked multi-cache strategy aimed at lowering memory latency and external memory bandwidth requirements and thus demonstrating the efficiency of 3D stacking to rethink architectures for obtaining unequalled performances in power efficiency.

I. INTRODUCTION

3D technologies represent a key enabler to new power-efficient memory architectures. In [8], authors demonstrate a reduction by a factor of two of the power consumption by stacking the DRAM memory directly on top of the processing circuit. Thanks to multiple wide-data access ports, the circuit I/O drive is reduced and wire length of the memory bus is globally shortened. In case of low power many-core architectures, stacking the DRAM memory on top of the processing tier also divides the power consumption by a factor of two to three while keeping the same performances [11]. The JEDEC low power memory subcommittee [1] recognizes the 3D technology promises in terms of bandwidth, power efficiency, scalability and tolerance to variability. The WIDEIO standard has been published to leverage this technology with products. In section II we will discuss the power efficiency advantages of this new memory interface standard. Our analysis is built on a physical implementation of a 3D circuit including a WIDEIO memory tier and an MPSoC tier.

WIDEIO and coarse grain partitioning are the first steps using 3D as an advanced packaging solution [12]. However, for more memory demanding applications, an external memory is still mandatory. In such cases, the high-density promises of Through-Silicon-Vias (TSV) can lead to innovative 3D cache architectures combining big capacities, high bandwidth and low power. In [5], authors show that the energy-delay product of a 130nm 3D SRAM is reduced by up to 80% in comparison to a conventional 2D SRAM. However, the heat effect of 3D stacking on SRAM leakage is not taken into account in this

study. When taking it into account, the energy-delay product of a four tiers SRAM 3D cache remains twice better than a conventional 2D cache in 25nm CMOS technology [6]. The same order of gain for a 4 layers 65nm 3D cache has been reported in [7]. Additionally the energy-delay product of a 96-entry processor register file can be reduced by 72% by using more signals for access ports thanks to wide vertical connections. Moreover, the power consumption of a L1 cache can be reduced by 75% compared to a conventional approach [10] when using innovative wide and adaptive vertical connections. Finally, by splitting logic functions of an Intel® Pentium® 4 processor in two dies, authors of [9] show that the total power consumption can be divided by two at constant performances.

Even if wide vertical interfaces provide a high bandwidth together with high power efficiency, TSVs negatively impact the integration density. To reduce the number of TSVs and improve the stacking flexibility, Networks-on-Chip (NoC) originally proposed for 2D many-core architectures [13][25] have been extended to the third dimension [14][15][18][19]. As the 3D clock distribution is a tough problem [16], asynchronous NoCs [17] can reduce the complexity of 3D integration. Moreover, the router can be made hierarchical to optimize its throughput [18] and the data can be serialized [19] to further reduce the number of TSVs. In [19], authors show a vertical link physical implementation providing 2 Gflits/s in a 65nm Low Power CMOS.

The high-density promise of TSV provides opportunities to innovative memory architecture partitioning. One hope is to find new solutions to minimize the huge memory bandwidth requirement which can stop the integration of more cores on a single die. We are addressing this issue through the proposal of a power-efficient 3D cache structure presented in section III.

II. WIDEIO STACKING

A. Memory Interface Evolution Towards 3D Integration

For embedded computing, new applications such as augmented reality lead to larger memory bandwidth and larger computation needs. While higher performance is achieved through a continuous increase in the number of processing elements, low power capability and high level of integration are driving the need for specific memory solutions and interfaces. Users and suppliers are collaborating to develop JEDEC standards [1] targeting those solutions by adapting existing interface (e.g. DDR3) to achieve lower power and lower cost. The newly introduced LPDDR3 mobile memory is based on a 32-bit data bus running up to 800MHz on a Dual

Data Rate (DDR) clock (Figure 1). The interface achieves 6.4GB/s data transfer per 32-bit memory channel. It is scheduled to improve the bandwidth of this interface by duplicating the number of channels and by improving the bit transfer rate with high-speed and low-swing IOs. The counterpart of this evolution is a high area cost for the adaptation layer (PHY), a high latency due to the additional low-swing protocol, a high design cost, and finally a complex validation on the application board.




Figure 1. Memory link features and associated peak bandwidth

An alternative technology based on 3D stacking is proposed by JEDEC with the WIDEIO standard published in December 2011. The interface with the SoC is based on a 4-channel bus of 128 bits each running up to 200MHz on a Single Data Rate (SDR) clock (Figure 1). It provides 12.8GB/s data transfer capability with the memory device. On package integration side (Figure 2), the WIDEIO memory is stacked on top of the processor SoC. The connections are located in the centre area of the memory die and consist of an array of TSVs and micro-bumps whose pitch is in the range of 40µm to 50µm.




Figure 2. LPDDR3 versus WIDEIO package integration

Compared with Package-on-Package (PoP) integration currently used for LPDDR3 (Figure 2), the power efficiency of the WIDEIO integration comes from the lower capacitance of the 3D interconnect and the high bandwidth capability from the wider data bus. Moreover, the low frequency (200 MHz) SDR type of data interface significantly simplifies the design of the pad interface. Signal-integrity issues are also singularly minimized with the short interconnect between dies. All these

features are facilitating the product bring-up on board and shortening its introduction.

B. Multi-Channel Memory Integration in MPSoC

The WIOMING MPSoC [2] has been designed, manufactured and tested by a collaboration of CEA-Leti, STEricsson and STMicroelectronics for demonstrating this new technology. The MPSoC backbone architecture (Figure 3) is organized around a 16-router Asynchronous NoC on which programmable units (Operator) as well as data management units (Data & Config.) are connected. This tile based architecture allows a seamless integration of the multichannel WIDEIO memory. It is done through four dedicated data management tiles as depicted in Figure 3 (WIDEIO Traffic Ctrl). Moreover, the WIDEIO traffic controller accommodates with parallel accesses and streaming transfers between tiles connected onto the asynchronous NoC. This unit is replicated four times as fully independent hard macros to fit with the logical and physical 3D interface footprint of the quad channel WIDEIO memory. In addition to its legacy task of data manipulation among the MPSoC processing operators, the WIDEIO traffic controller unit offers peak-rate data flow per DRAM channel. It handles, in parallel, high-bandwidth data read/write between on-chip SRAM banks and WIDEIO memory channel and low-latency medium-bandwidth data manipulation between NoC and WIDEIO.

This architecture approach brings a real performance improvement when dealing with stream-based application.

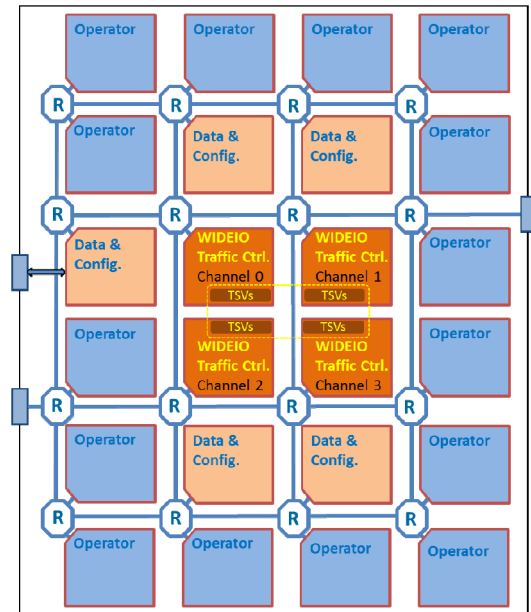


Figure 3. WIOMING tile-based architecture

C. WIDEIO PHY Adaptation Layer

On physical implementation side [4], the PHY layer which consists is providing DRAM timing compliant signal, has been placed in a regular way close to the TSV matrix. The PHY consists in a 128-bit wide DRAM data capture interface that is

able to operate up to 200 MHz in SDR mode. In our architecture, a DLL-less fixed quarter cycle delayed clock has been implemented to avoid set-up time violation during data capture at high frequency. This WIDEIO PHY solution is less constrained than traditional architecture designed for LPDDR3 running up to 800MHz. It permits simple design in favor of power efficiency thanks to less sensitivity to variability and less buffering needs.

A validation of the PHY has been made on an application board where heavy data transfers between the SoC and the memory have been implemented. To operate at full speed, we've used incremental burst to get one data access per cycle. By checking the data integrity, the application tests have shown stable operations of our power optimized PHY as predicted by our simulations.

D. WIDEIO Die-to-Die Interconnect

WIDEIO Mobile DRAM which uses chip-level 3D stacking with Through Silicon Via (TSV) interconnects and memory chips directly stacked on the MPSoC requires innovative design solutions for vertical signal propagation. To ensure the integrity and robustness of 3D signals through the heterogeneous die stack, micro-buffers (Figure 4) have been placed between the TSVs as close as possible from each input/output (IO) signal.




Figure 4. WIOMING vertical connections

Compared with PoP package, the connections between dies are greatly shortened with 3D integration. In our design case, the wire length corresponds to the die thickness which is 80µm. The capacitive load of such short interconnect is reduced and does not require power hungry signal buffers for off-package propagation. The power consumption of the IO interconnect is consequently drastically reduced compared with traditional memory interfaces.

E. DFT and Power Measurements

Classical Design-for-Test (DFT) techniques such as TAP based scan has been applied to 3D signals for connection

failure detection. In addition, some 3D specific DFT techniques are required like the SoC embedded Memory Built-in-Self-Test (MBIST). This test is intended to analyze the possible memory failure after stack assembly. All DFT features are accessible from the SoC JTAG port and managed by an IEEE1500 TAP controller. We believe that the MBIST test embedded in the SoC and run at speed is a good candidate to be use in an Automated Test Equipment (ATE) environment for precise 3D stack power measurement. During the peak activity section of this test pattern, we have measured on the ATE the power consumption of the complete 3D stack at ambient temperature (25°C). The total power measured on the 3D-IC considering peak activity on the SoC Traffic controller, the 3D IO link (12.8 GB/s) and the WIDEIO memory plane is 283mW. It represents a power efficiency of 2.8 pJ/bit for the full WIDEIO system. If we remove the SoC part of the power consumption to only consider the IO link and memory, the power consumption becomes 170 mW. It represents less than half of the consumption of a LPDDR3 memory with IO load corresponding to board-level integration (406 mW in [24]).

3D stacking is a unique opportunity enabling memory-interconnect evolution to higher bandwidth with a limited power impact thanks to very short connections between dies. WIDEIO prototypes have been already designed and validated [3][20] making the 3D current technology ready for production of power efficient memory interfaces. Next generation of 3D integration technology is under development and will bring architectural opportunities with power efficient partitioning thanks to higher TSV density. Next section explains how we can achieve a new paradigm in power efficiency with advanced cache architecture built around next generation of 3D integration technology.

III. HIGH BANDWIDTH LOW POWER 3D STACKED CACHE

In many-core architectures, the scaling of the number of cores is limited by the bandwidth of the external memory [21], it is the memory wall problem. Embedding the main memory on top of the processing architecture is not possible in a High Performance Computing (HPC) many-core context. HPC applications indeed need tens of gigabytes of memory which is too much for 3D integration. WIDEIO for example can stack no more than 8 GB of memory. Moreover the memory density increase will not fill this gap because in the same time applications will need more memory due to the multiplication of cores.

So, a classical approach is to use cache hierarchies to reduce the memory pressure and thus minimize the memory wall problem: instead of stacking all the memory on top of processors, we add a big 3D distributed cache. Stacking fine-grained caches interconnected with tens of 3D NoC links as shown in Figure 5 can satisfy the bandwidth and cache size requirements along. These caches behave as big non-uniform cache architecture (NUCA). The possibility of stacking fine-grained caches on top of the processors opens new perspectives on the many-core architecture and allows developing new efficient memory architectures.

A. Distributed 3D Cache Architecture

Figure 5 depicts the 3D cache architecture stacked over the many-core architecture. The bottom tier is the many-core while on top there are the 3D cache tiers. The 3D cache tier architecture is fully stackable and the number of cache tiers can be chosen at assembly time. So, the total cache size can be modified according to application’s needs late in the design process. Finally, the inter-tier links use the high density TSVs of recent 3D technology in order to minimize its area overhead. Therefore, very large vertical interconnections can provide a high vertical bandwidth to the distributed 3D cache architecture

A 3D cache tier is a mesh topology of cache tiles interconnected through a 3D version of the DSPIN [23] NoC. A 3D NUCA cache meeting these requirements has been proposed in [22]. The 3D cache architecture is composed of cache access controllers and cache tiles. The cache access controllers are memory mapped on the global memory space while the cache tile is a fully functional cache. Therefore, when a processing unit (PU) sends a request, the cache access controller dispatches the request to a cache tile. The 3D NoC links are 64 bits wide and the NoC cycle time is 1 GHz.

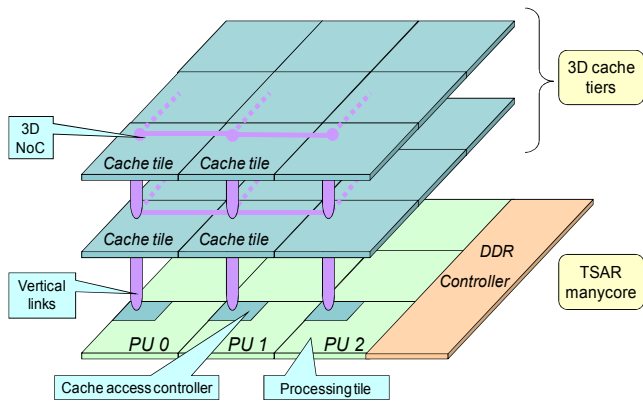


Figure 5. 3D adaptive cache

In this architecture, the cache access controllers located in the bottom tier dispatch the requests from processing units to cache tiles by implementing software controlled data placement functions described in [22]. Therefore the limited amount of cache can be optimized to speedup a specific application. The cache access controllers implement a dedicated configuration interface controlled by the operating system (OS).

A cache tiles is a fully functional and autonomous cache. It is composed of a cache bank, a cache directory, the cache logic and the 3D NoC routers. It is able to handle read and write requests from multiple initiators in a fully distributed manner. In case of cache MISS, the cache tile generates a request to the external memory controller for cache line replacement. It also implements additional capabilities, like selective cache invalidation for cache coherency and power management.

Architectural low-power techniques were used on the design of the cache tile. For example, it uses a very wide cache data path in order to reduce its clock frequency. Therefore, high

density and low power RAM memories are well suited for this architecture.

B. Performances of the 3D Cache

A four tiers 3D cache with a mesh topology of 8x8 tiles per tier (256 total tiles), experimentations have shown that this 3D cache can provide a bandwidth of up to 745 GB/s. This is about sixty times more than the current WIDEIO bandwidth. It enables the scaling up of many-core architectures without reaching the memory wall issue.

The software controlled data placement functions allows reallocating the cache resources to the applications. Therefore, up to 50% reduction of both execution time and traffic to the external memory can be achieved to favor a particular application. Finally, in case of heterogeneous workloads, such cache configuration schemes optimize the exploitation of the limited memory bandwidth resource.

C. Physical Implementation of a Cache Tile

A hardware implementation of a 1 MB cache tile has been done in the 28nm STMicroelectronics CMOS LP process. We used 5µm-wide TSVs with a 10µm pitch and the data of the cache were stored in SRAM memories. The cache tile footprint is a square of 1.5mm². The achieved useful memory density is as high as 77% of the tile area, which validates our fine-grained 3D approach.

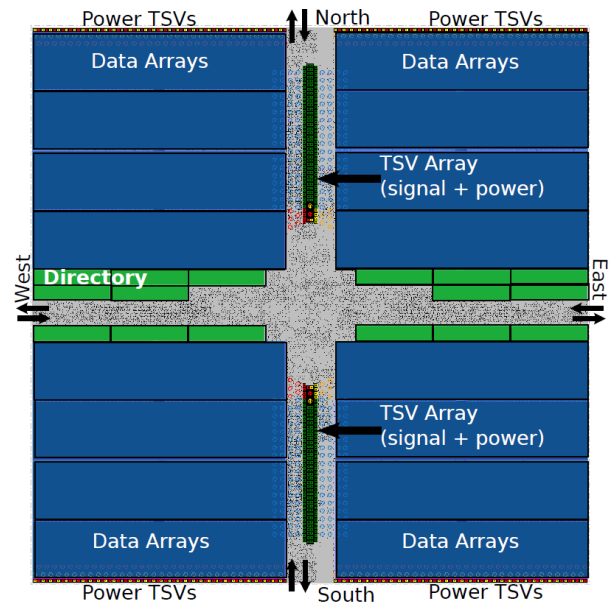


Figure 6. Backend of a 1MB 3D cache tile

The tile includes 284 signal TSVs and 280 power TSVs. The number of power TSVs has been computed from WIDEIO power distribution infrastructure. Thus, the 5µm-wide TSVs represent only 4% of the total area, but 10µm TSVs (40x50µm of pitch) would have taken up to 33% of the tile area. So it shows that, for such fine-grained architectures, high density TSVs are mandatory to achieve adequate memory density.

Otherwise, the cache granularity should be coarser (i.e. embed more memory) in order to hide the area overhead of the TSVs.

The copper pillars have a 20 μ m pitch which is bigger than the pitch of TSVs. So, to build compact TSVs arrays, back side redistribution layer (RDL) routing is needed between TSVs and back bumps. The Figure 7 shows a part of the TSVs and back bumps arrays with back side tracks.

Finally, the use of eDRAM memories is an interesting solution to increase the cache density and thus postpone the memory wall. However, its power consumption is increased compared to SRAM memories. Moreover, eDRAM memories suffer more from variability and power consumption against temperature variations, which can be a challenge to deal with in a stacked many-core architecture.

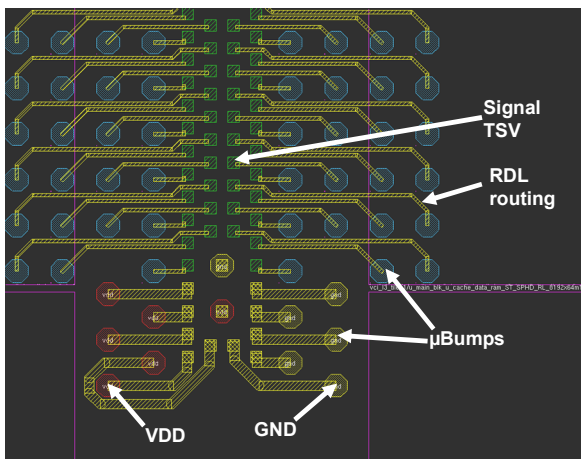


Figure 7. TSV and back bump arrays

D. Power Consumption

The power consumption of the cache tile has been computed in post-backend simulations using synthetic and real traffic patterns. The power consumption was computed under typical PVT corner (typical operating conditions, 1.0V and 25°C). The capacitive load of the high density TSVs and bumps has been modeled as a 50fF load, which is four times smaller than the capacitive load of a WIDEIO stack. The synthetic traffic pattern is made up of cache request with 100% cache hit rate and a 50% reads and 50% writes of 64 bytes fixed length random access patterns. Under this workload, the cache tile provides a maximum throughput of 56.5 Gb/s for 100000 requests.

TABLE I. POWER CONSUMPTION OF A CACHE TILES UNDER SYNTHETIC TRAFFIC PATTERNS

Idle power (incl. clocks)	18.4 mW
Peak power	< 170 mW
Read hit energy (100% read)	1.31 pJ/bit
Write hit energy (100% write)	3.41 pJ/bit
Power efficiency ¹	2.31 pJ/bit

(1) For the maximum bandwidth measured with a 50% reads and 50% writes of 64 bytes length traffic

Table I summarizes the power consumption of a single 1 MB cache tile using the above-mentioned synthetic traffic. The power efficiency is 2.3 pJ/bit, which is 18 % better than the WIDEIO efficiency under the same operating conditions.

Each power TSV pair (VDD+ground) are expected to feed a current of 5 mA in peak, giving a total of 700 mA per tile. As the peak power consumption of the cache tile is inferior to 170 mW (i.e. 170 mA at 1.0 V), up to four tiers of 3D cache can be stacked without exceeding the capacity of power TSVs.

Finally, we have computed the power consumption of a single 3D cache tier of 4x4 cache tiles using real traffic patterns. A cycle and bit accurate SystemC model of the 16 tiles/64 cores TSAR [25] many-core were used to execute Splash2 [26] HPC applications. The 3D cache serves here as a distributed L3 cache made up of 16 tiles of 1 MB per tile. The size of the cache tile is chosen to be equal to the size of a processing tile. The Table II shows the average power consumption on multiple applications and stacking configurations. The minimum and maximum power consumption is computed on a millisecond integration period basis. This table shows that the total 3D cache peak power consumption doesn't exceed 800mW for 1-tier configuration and 2W for 4-tier stack when running compute intensive HPC applications.

TABLE II. TOTAL POWER CONSUMPTION OF DIFFERENT 3D CACHE CONFIGURATIONS UNDER SPLASH TRAFFIC PATTERNS

Application	3D cache	Total power consumption min/avg/max (mW)
FFT 218 elements	1 tier 16 MB	298.1 / 335.2 / 779.1 (82 ms execution time)
Ocean 256x256 grid	1 tier 16 MB	298.1 / 333.6 / 660.1 (135 ms execution time)
LU 1024x1024 matrix	1 tier 16 MB	298.1 / 306.1 / 354.3 (458 ms execution time)
FFT 220 elements	4 tiers 64 MB	1192 / 1242.4 / 1953.7 (290 ms execution time)

So, this 3D cache architecture delivers a very high throughput with contained power consumption. Thanks to a low power architecture and CMOS process, it reaches a better power efficiency than WIDEIO, trading the quantity of memory for a higher bandwidth. Furthermore, the number of power TSVs is enough to stack up to four tiers of 3D cache, thus embedding up to 256 MB of memory on top of a 256 cores processing tier.

IV. CONCLUSION

In this paper we have shown two applications for 3D stacking technology leading to power-efficient memory architecture solutions. The first one is focused on solving the power hungry memory interface issue when targeting ultra-high bandwidth especially for embedded applications. We have demonstrated, on a WIDEIO silicon prototype, a drastic power efficiency improvement by stacking the memory on top of the

MPSoC compared with classical off-package memory interface solutions. The second memory architecture solution shown in this paper is focused on solving the latency and bandwidth issue for advanced cache architecture in many-core context. With simulations done on a hardware implementation, we demonstrate that next generation of high-density TSV leverages new power-efficient cache architecture required in HPC and data-servers.

REFERENCES

- [1] JEDEC web site: <http://www.jedec.org>
- [2] F. Clermidy, D. Darve, D. Dutoit, W. Lafi, P. Vivet, "3D Embedded multi-core: Some perspectives," *DATE 2011*, pp.1-6, March 2011.
- [3] <http://blog.stericsson.com> – July 2012
- [4] P. Vivet, V. Guerin, "A Three-Layers 3D-IC Stack including Wide-IO and 3D NoC – Practical Design Perspective", *Presentation at the 2011 RTI 3D ASIP conference, San Francisco, USA*, Dec 2011.
- [5] S. S. Wong et A. El Gamal, "The prospect of 3D-IC," in *Custom Integrated Circuits Conference, 2009. CICC'09. IEEE*, 2009, p. 445–448.
- [6] Yuh-Fang Tsai, Feng Wang, Yuan Xie, N. Vijaykrishnan, et M. J. Irwin, "Design Space Exploration for 3-D Cache," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 16, n° 4, p. 444-455, avr. 2008.
- [7] K. Puttaswamy et G. H. Loh, "3D-Integrated SRAM Components For High-Performance Microprocessors," *Computers, IEEE Transactions on*, vol. 58, n° 10, p. 1369–1381, 2009.
- [8] M. Facchini, T. Carlson, A. Vignon, M. Palkovic, F. Catthoor, W. Dehaene, L. Benini, et P. Marchal, "System-level power/performance evaluation of 3D stacked DRAMs for mobile applications," in *Proceedings of the Conference on Design, Automation and Test in Europe*, 2009, p. 923–928.
- [9] B. Black, M. Annavaram, N. Brekelbaum, J. DeVale, L. Jiang, G. H. Loh, D. McCaule, P. Morrow, D. W. Nelson, D. Pantuso, et others, "Die stacking (3D) microarchitecture," in *Microarchitecture, 2006. MICRO-39. 39th Annual IEEE/ACM International Symposium on*, 2006, p. 469–479.
- [10] T. Ono, K. Inoue, et K. Murakami, "Adaptive cache-line size management on 3D integrated microprocessors," in *SoC Design Conference (ISOCC), 2009 International*, 2009, p. 472–475.
- [11] T. Kgil, S. D'Souza, A. Saidi, N. Binkert, R. Dreslinski, T. Mudge, S. Reinhardt, et K. Flautner, « PicoServer: using 3D stacking technology to enable a compact energy efficient chip multiprocessor », *SIGOPS Oper. Syst. Rev.*, vol. 40, no 5, p. 117–128, oct. 2006.
- [12] J.U. Knickerbocker et al., "2.5D and 3D technology challenges and test vehicle demonstrations," *Electronic Components and Technology Conference (ECTC), 2012 IEEE 62nd*, pp.1068-1076, June 2012
- [13] D. Melpignano, L. Benini, E. Flaman, B. Jago, T. Lepley, G. Haugou, F. Clermidy, D. Dutoit, "Platform 2012, a many-core computing accelerator for embedded SoCs: Performance evaluation of visual analytics applications", *Design Automation Conference*, June 2012.
- [14] A. Sheibanyrad, F. Petrot, A. Jantsch, "3D Integration for NoC-Based SoC Architectures", *Springer 2011*, ISBN 978-1-4419-7617-8.
- [15] W. Lafi, D. Lattard and A. Jerraya "A Stackable LTE Chip for Cost-Effective 3D Systems", in *IPSSJ Transactions on System LSI Design Methodology*, Vol 5, 21 fev 2012.
- [16] V. F. Pavlidis, I. Savidis, and E. G. Friedman "Clock Distribution Networks for 3-D Integrated Circuits", *Proc. of IEEE Custom Integrated Circuits Conference, CICC'08*, 2008.
- [17] Y. Thonnart, P. Vivet and F. Clermidy, "A Fully-Asynchronous Low-Power Framework for GALS NoC Integration", *Proc. of Design And Test in Europe, DATE'10*, Dresden, Germany, March 2010.
- [18] W. Lafi, D. Lattard, A. Jerraya, "An Efficient Hierarchical router for 3D NoC Architecture", *proc. of IEEE International Symposium on Rapid System Prototyping*, RSP'2010, Fairfax, USA, June 2010.
- [19] F. Darve, A. Sheibanyrad, P. Vivet, F. Petrot, "Physical Implementation of an Asynchronous 3D-NoC Router using Serial Vertical Links", *IEEE ISVLSI'2011*, Chennai, India, July 2011.
- [20] J.-S. Kim et al., "A 1.2v 12.8gb/s 2gb mobile wide-i/o dram with 4x128 i/os using tsv-based stacking," in *ISSCC 2011, feb. 2011*, pp. 496–498.
- [21] B. M. Rogers, A. Krishna, G. B. Bell, K. Vu, X. Jiang et Y. Solihin, "Scaling the bandwidth wall : challenges in and avenues for CMP scaling". *SIGARCH Comput. Archit. News*, 37(3):371–382, juin 2009.
- [22] E. Guthmuller, I. Miro-Panades, and A. Greiner, "Adaptive stackable 3d cache architecture for manycores," in *VLSI (ISVLSI), 2012 IEEE Computer Society Annual Symposium on*, aug. 2012, pp. 39–44.
- [23] I. Miro-Panades, A. Greiner, and A. Sheibanyrad, "A Low Cost Network-on-Chip with Guaranteed Service Well Suited to the GALS Approach," *Nano-Networks and Workshops*, 2006.
- [24] Yong-Cheol Bae; Joon-Young Park; Sang Jae Rhee; Seung Bum Ko; Yonggwon Jeong; Kwang-Sook Noh; Younghoon Son; Jaeyoun Youn; Yonggyu Chu; Hyunyoong Cho; Mijo Kim; Daesik Yim; Hyo-Chang Kim; Sang-Hoon Jung; Hye-In Choi; Sungmin Yim; Jung-Bae Lee; Joo Sun Choi; Kyungseok Oh; , "A 1.2V 30nm 1.6Gb/s/pin 4Gb LPDDR3 SDRAM with input skew calibration and enhanced control scheme," *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pp.44-46, 19-23 Feb. 2012.
- [25] "TSAR project website." [Online]. Available: <https://www-soc.lip6.fr/trac/tsar>
- [26] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, and A. Gupta, "Splash-2 programs: Characterization and methodological considerations," in *Conference Proc. - Annual Int. Symp. on Computer Architecture, ISCA, 1995*, pp. 24–36.