Thermal Modelling of 3D Stacked DRAM with Virtual Platforms

Matthias Jung*,
MohammadSadegh Sadri*,
Norbert Wehn*

* *Microelectronic System Design Research Group University of Kaiserslautern*
Erwin-Schroedinger-Strasse 67663 Kaiserslautern, Germany

ABSTRACT

Heterogeneous 3D integrated systems with Wide-I/O DRAMs are a promising solution to squeeze more functionality and storage bits into an ever-decreasing volume. Unfortunately, with 3D stacking the challenges of high power densities and thermal dissipation are exacerbated. In order to provide a framework for exhaustive design space explorations, we developed an advanced Transaction Level (TLM) based virtual platform that models the performance, power and thermal behaviour of a 3D-integrated MPSoC with Wide-I/O DRAMs in detail.

KEYWORDS: 3D-DRAM, Wide-I/O, Refresh, Thermal, Simulation, TLM2.0

1 Introduction

Energy and thermal dissipation are limiting the efficiency of today's green computing solutions. More than 40% of the system energy is consumed by DRAMs [LZZZ13] in existing platforms. 3D packaging of systems starts to break down the memory and bandwidth walls. However, this comes at the price of increased power density and less horizontal heat removal capability of the thinned dies. The thermal issues of 3D ICs cannot be solved by tweaking the technology and circuits alone.

In fact, a 3D stacked SoC aggravates the thermal crisis and forces enhancements in the architecture and memory organization. Early detection of architectural shortcomings and thermal hazards is crucial to the design of sub-20-nm 3D chips. For instance, current microprocessor architectures are inefficient for running datacenter workloads mainly because of the mismatch between the workload characteristics and the organization of the memory subsystem [F+12]. Consequently, the detailed analysis of the memory subsystem is very important as it unveils possible bottlenecks and issues which impact the system energy and efficiency.

Therefore, we perform an in-depth study on performance, timing, power and leakage of Wide-I/O DRAMs and track its key parameters with temperature change. Through care-

[1]E-mail: {jungma,sadri,wehn}@eit.uni-kl.de

ful evaluation of temperature distributions in a 3D IC with Wide-I/O DRAMs, we propose architectural enhancements for the DRAM subsystem as well as thermal management solutions which improve energy consumption. We consider the lateral and vertical variation in temperature of the 3D DRAM dies and refresh each of the DRAM banks at a separate rate according to its own temperature. In order to assess and quantify the advantages of our proposed ideas, we build a suitable virtual infrastructure which considers all key characteristics of a 3D MPSoC with Wide-I/O DRAMs in detail.

2 Virtual Platform

The analysis of the memory subsystem requires a timing accurate behavior of the CPU cores. Therefore, the *gem5* [BBB+11] architecture level full-system simulator is selected for this purpose, since it models system operations in detail and generates realistic traces of memory accesses, which can be replayed very fast inside the TLM environment [JSW13].

Figure 1 shows our developed simulation infrastructure [SJW+14] that consists of fast and flexible TLM models of the Wide-I/O DRAM subsystem [JWWC13, JWB+13][2] as well as gem5 trace players, which replay pre-recorded memory accesses (after L2 cache) from the gem5 simulator, running different benchmarks on Android.

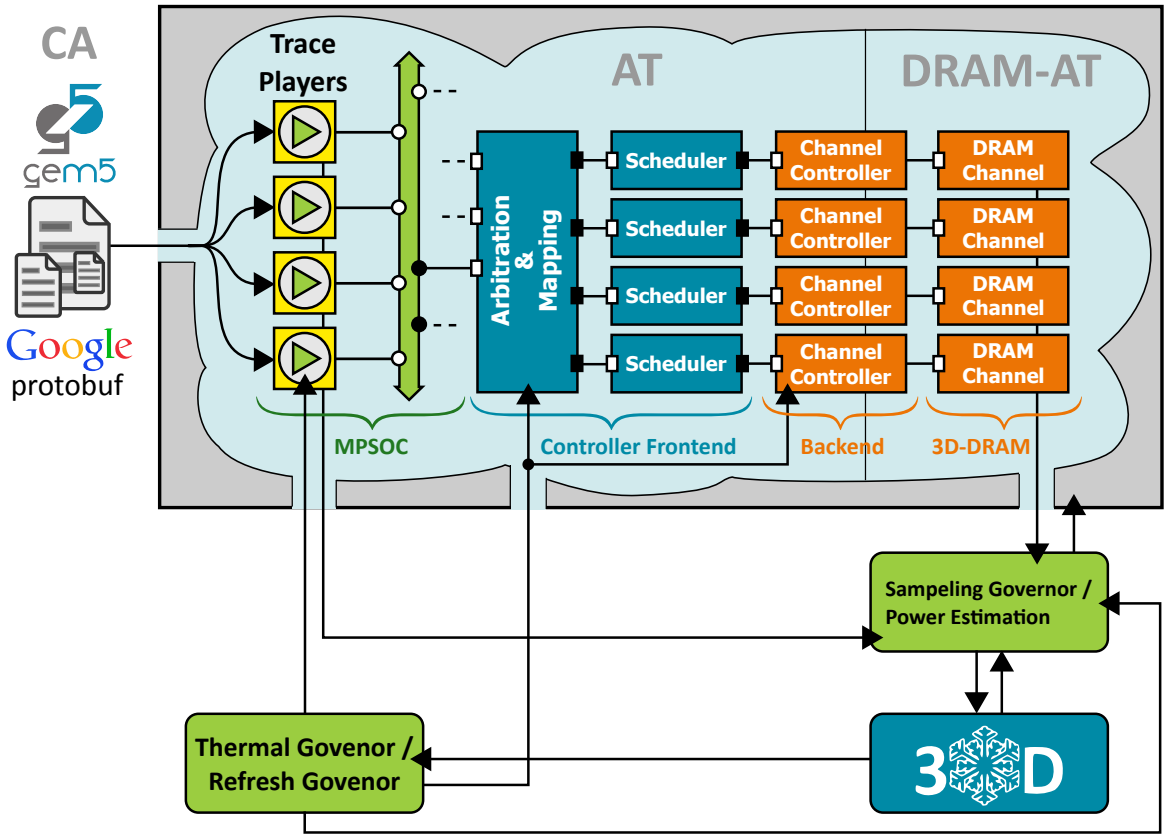

Figure 1: Simulation Environment

[2]The DRAMSys models will be available on http://www.uni-kl.de/3d-dram

The performance statistics of CPU cores and DRAMs are used to estimate their power, which is used as an input for the thermal co-simulation. We build this thermal co-simulation using the 3D-ICE thermal simulator [SVR+], which supports definition of layers with non-homogeneous materials in the chip stack. The results of the thermal simulation are used to govern the refresh rate of the DRAM and the voltage and frequency of the cores. This sophisticated closed loop, enables us to explore new controller architectures and techniques to optimise 3D-stacked memory systems in terms of performance, power, and thermal impacts.

3 Temperature Variation Aware Bank-Wise Refresh

The mentioned lateral and vertical temperature variations cause significant differences in the required refresh rate of each DRAM bank. Due to these observations, we implemented the following key idea: instead of defining the refresh rate based on the maximum temperature seen across the entire channel and refreshing all DRAM banks at the same rate, we select the refresh rate of each bank separately based on its own maximum temperature (see Figure 2a). We have extended our DRAM subsystem model to support handling of separate perbank refresh commands. This increases the overall refresh period (makes refreshes happen less frequently) and improves the energy consumption.

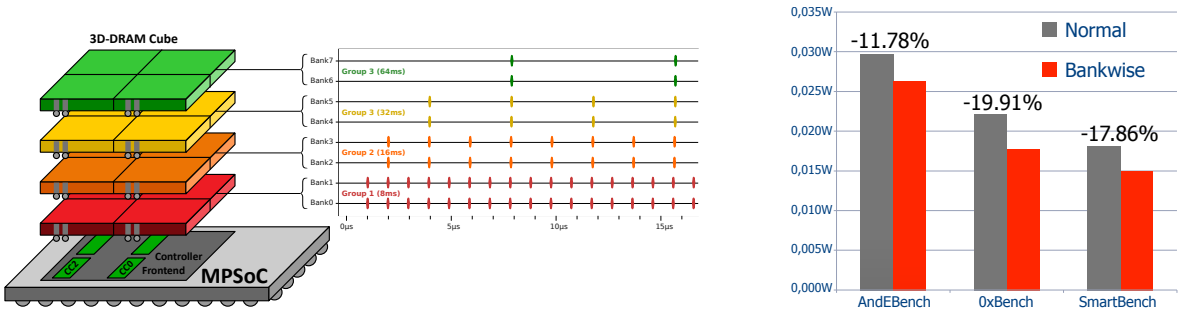

(a) (b)

Figure 2: Concept and Results of Bank-Wise Refresh

To show the advantage of this idea we executed two sets of simulations; In the first set the bank-wise refresh is disabled. For each DRAM channel, the refresh governor finds the maximum temperature of the channel. Then suitable refresh period will be selected and used for all DRAM banks of the channel. In the second set, the bank-wise refresh is active and different refresh rates are assigned to each bank group, as shown in Figure 2a. We run Android 2.3 Gingerbread on the virtual platform and use three real-world well-known benchmarks to stress the CPUs and the memory. Our selected benchmarks are: *AndEBench, 0xBench, and SmartBench*.

Figure 2b shows that the bank-wise refresh approach can save up to $\sim 20\%$ of refresh power. For all typical workloads we see an average improvement of $\sim 16\%$. In near future DRAMs, half of the DRAM power will be related to refresh [LJVM]. Thus, the proposed idea will significantly improve the total energy consumption.

References

[BBB+11] N. Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood. The gem5 Simulator. *SIGARCH Comput. Archit. News*, 39, 2011.

[F+12] Michael Ferdman et al. Clearing the Clouds: a Study of Emerging Scale-Out Workloads on Modern Hardware. *SIGARCH Comput. Archit. News*, 40(1):37–48, March 2012.

[JSW13] Matthias Jung, MohammadSadegh Sadri, and Norbert Wehn. Virtual platforms for fast memory subsystem exploration using gem5 and tlm2.0. In *Poser Abstracts of Ninth International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems (ACACES13)*, volume 9, pages 153–156. Academia Press, Ghent, Belgium, 2013.

[JWB+13] Matthias Jung, Christian Weis, Patrick Bertram, Gunnar Braun, and Norbert Wehn. Power Modelling of 3D-Stacked Memories with TLM2.0 based Virtual Platforms. In *Synopsys Usergroup Conference (SNUG)*. Synopsys, 2013.

[JWWC13] Matthias Jung, Christian Weis, Norbert Wehn, and Karthik Chandrasekar. TLM modelling of 3D stacked wide I/O DRAM subsystems: a virtual platform for memory controller design space exploration. In *Workshop on Rapid Simulation and Performance Evaluation 2013 (RAPIDO) HiPEAC Conference 2013*, 2013.

[LJVM] J. Liu, B. Jaiyen, R. Veras, and O. Mutlu. RAIDR: Retention-Aware Intelligent DRAM Refresh. In *Proc. of ISCA 2012*.

[LZZZ13] J. Lin, Hongzhong Zheng, Zhichun Zhu, and Zhao Zhang. Thermal Modeling and Management of DRAM Systems. *IEEE Trans. on Computers.*, 2013.

[SJW+14] Mohammadsadegh Sadri, Matthias Jung, Christian Weis, Norbert Wehn, and Luca Benini. Energy optimization in 3D MPSoCs with wide-I/O DRAM using temperature variation aware bank-wise refresh. In *Proceedings of the conference on Design, Automation & Test in Europe*, page 281. European Design and Automation Association, 2014.

[SVR+] A. Sridhar, A. Vincenzi, M. Ruggiero, Thomas Brunschwiler, and D. Atienza. 3D-ICE: Fast compact transient thermal modeling for 3D ICs with inter-tier liquid cooling. In *Proc. of ICCAD 2010*.