

Architecture Simulation Framework for 3D IC

William J. Song Saibal Mukhopadhyay Sudhakar Yalamanchili
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA 30332
wjhsong@gatech.edu, {saibal, sudha}@ece.gatech.edu

ABSTRACT

The practical limitation of implementing 3-dimensional integrated circuits is the increased thermal stress. The integration of liquid cooling with 3D IC can solve the thermal problems, allowing up to $1,000\text{W}/\text{cm}^2$ of power density [1, 2]. Thus, the significant performance improvement can be made through 3D ICs and liquid cooling integration. However, increased transistor and power density will address failure concerns, threatening the lifetime reliability of microprocessors [5]. Therefore, 3D ICs must be designed under the constraints of microarchitecture, energy, thermal, area, and reliability. There had been many efforts to develop accurate and fast modeling of these 3D IC physical properties. We note that the simulation models of these physical phenomena must be designed to interact with each other to correctly capture the tradeoffs among 3D IC design constraints. In this paper, we present the methodology for 3D IC simulations and present a novel simulation framework for the 3D IC design space exploration.

1. INTRODUCTION

Aggressive technology scaling faces the power and thermal limitations, entering the dark silicon era. A number of dynamic execution techniques (i.e., computational sprinting, specialization, voltage-frequency scaling) are being proposed to overcome the dark silicon issue [13, 14, 15], but they do not provide the essential breakthrough solution to overcome the imminent physical limitations. 3-dimensional integrated circuits (3D ICs) had been emerging as a promising solution. 3D ICs vertically stack multiple layers of processing units (i.e., core, cache, memory, interconnection network) in a processor package, significantly reducing the distance between the units and thus improving the performance [1, 2, 3]. However, the major drawback of the 3D IC is the increased power and thermal densities per area or volume, which makes the 3D integration less feasible with conventional air cooling whose cooling efficiency is typically limited around $100\text{-}150\text{W}/\text{cm}^2$ [1]. The introduction of advanced cooling techniques (i.e., liquid cooling) built in the 3D package can effectively solve the thermal issue. The liquid cooling compared with air cooling has significantly higher cooling efficiency, which makes the processor to endure much larger power density and operate at higher voltage and frequency levels

[1]. In the mean time, the increased power density as well as transistor density per volume address the lifetime reliability concerns. In this regards, the full-system simulation capability is essential for the design space exploration of 3D ICs, including microarchitecture, power, temperature, cooling, and reliability models and most importantly the interactions among the physical phenomena (e.g., power-thermal, degradation-power interactions).

In this paper, we first review the physical models for the 3D IC design space exploration and present the devised simulation framework for interactive microprocessor physics.

2. MODELING PHYSICAL PROPERTIES

This section describes the modeling methodology for simulating 3D IC microarchitecture and physical phenomena.

2.1 Microarchitecture and Energy Models

Energy dissipation is the fundamental source of microprocessor physics, resulted by computing. Energy (or power) dissipation in the processor is largely composed of two parts; leakage and dynamic dissipations [17, 18]. Leakage dissipation refers to the energy usage to keep the processor active without actual computing processes (i.e., idle state). The leakage energy is exponentially dependent on the temperature increase and proportional to the active time, and thus reducing the active-state temperature and power-gating help reduce the leakage dissipation [13]. On top of the leakage, the dynamic energy dissipation occurs as the result of actual computing. In microarchitecture simulations, the dynamic energy is computed at the functional block level, called *modules*. The definition and granularity of modules differ from what specific energy/power models are capable of estimating and how much detail is required. In general, the modules are defined as functional components of the microarchitecture such as data cache, ALU, instruction queue, etc. (see Figure 1) The modules are typically analyzed at the circuit-level, where architectural and technology parameters are used to estimate the dynamic and leakage energy based on the generic design of a microarchitecture component. The dynamic energy is computed by associating switching activity count with unit switching energy, where the activity count can be collected from architecture simulations. The activity counts indicate how many times an architectural component was accessed during the observation period, and thus high access counts imply large dynamic energy dissipation.

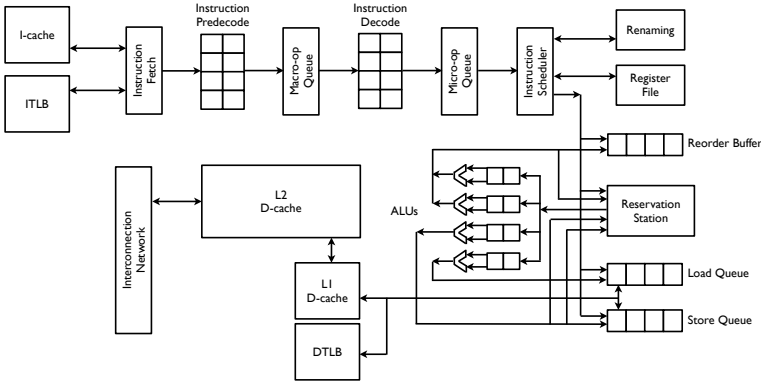


Figure 1: Microarchitecture design comprised of functional modules [16, 17].

2.2 Thermal and Cooling Models

The conventional air cooling is implemented by placing a heat sink on the processor package, and the heat is removed via the heat sink and fan. Figure 2 shows the air cooling model, where the 3D package is modeled as the thermal RC grid. The cooling efficiency of air cooling is generally around 100-150W/cm², which is insufficient for high performance 3D processors [1].




Figure 2: 3-dimensional thermal grid cell (left) and package design (right) of conventional air cooling [21].

The shift of cooling technique from using air to liquid can significantly improve the cooling efficiency and thus bring about the performance growth. The liquid cooling is implemented by fabricating the micro-channels inbetween the stacked layers of the 3D package and attaching micro-channels to the heat sink as shown in Figure 3. The coolant liquid (i.e., water) flows through the micro-channels to remove the heat away from the processor. The liquid cooling can remove the heat up to 1,000W/cm² which is significantly higher than the air cooling efficiency [1, 2, 3].




Figure 3: 3-dimensional thermal cell (left) and package design (right) of liquid cooling in 3D IC.

2.3 Reliability Models

Continued multicore scaling with shrinking feature size threatens the reliability of microprocessors. Srinivasan et al. in [5] showed that the mean-time-to-failure (MTTF) of processors is diminishing each technology along with voltage and frequency scalings. The presence of many cores in the same processor chip will inevitably experience the uneven degradation across the silicon die and between layers. In addition, architectural heterogeneity or specialization across and between silicon dies in the 3D IC will exacerbate the non-uniformity due to inherit difference in power and thermal features by executing different tasks through different microarchitecture designs. For instance, the separate configuration of core, cache, interconnection network, and memory layers will have different rates of degradation over the lifetime, and the mixture of multiple microarchitecture components (e.g., asymmetric or heterogeneous cores) will worsen the unevenness of the degradation. The chip-level lifetime will be eventually limited by the earliest failure component in the processor. Therefore, the reliability is the constraint that must be taken into account for the 3D IC design space exploration.

Several microprocessor failure models were presented in [4, 5, 6, 7, 8]. The following summarizes the known failure mechanisms of microprocessors:

- Electromigration (EM): A directional transport of electrons and metal atoms in interconnect wires leads to degradation and eventual failure.
- Time-dependent dielectric breakdown (TDDB): Wearout of gate oxide caused by continued application of electric field leads to electric short between gate oxide and substrate.
- Hot carrier injection (HCI): Electrons that capture sufficient kinetic energy overcome the barrier to gate oxide and cause the threshold voltage shift and degradation.
- Negative bias temperature instability (NBTI): The holes trapped from the gate cause the threshold voltage shift and timing error. The switching between negative and positive gate voltages cause degradation and recovery of the NBTI effect.
- Stress migration (SM): Mechanical stress due to the differences between the expansion rates of metals causes the failure.
- Thermal cycling (TC): Fatigue accumulates with temperature cycles with respect to the ambient temperature.

In the presence of possible failure risks, the microprocessor is operational only if none of above failure mechanisms ($r \in$ risks) occur over the time ($t \in [t_0, t_n]$) for all comprising components ($c \in$ component $[c_1, c_m]$) in the processor; the series system model [9]. Thus, the total failure probability of the processor can be computed as follows:

$$P_{total} = 1 - S_0 \prod_{c=1}^m \prod_{r \in \text{risks}} \prod_{i=1}^n \left\{ \frac{1 - P_{c,r}(t_i - t_{i-1})}{\lambda(T(i), f(i), V(i), a(i), g(i)))} \right\} \quad (1)$$

With dynamic operation of the processor, the failure probability λ is time-varying with temperature (T), clock frequency (f), supply voltage (V), switching activity factor (a), and power gating state ($g \in [\text{on}, \text{off}]$).

2.4 Interactive Simulation Interface

The physical characteristics presented in the earlier sections are not independent but strongly interactive phenomena. For example, the power dissipation of an architecture module leads to heat dissipation and temperature increase, which in turn increases the leakage power. In addition, the electrical and thermal stresses cause the device-level degradation which affect the timing, dynamic and leakage powers, and eventually alter the thermal results. Therefore, such interactions among physical phenomena must be correctly captured in the 3D IC architecture simulations, and we devised a novel simulation interface, called the *Energy Introspector* (EI), to support such capabilities.

The practical problem of designing framework is correlating the separate physical models (i.e., energy/power, thermal/cooling, and degradation/reliability). The EI connects the models via *pseudo components* which are the abstract representation of processor modeling components; *package*, *partition*, and *module*. The EI has an additional pseudo component type, *sensor*, to emulate the sensor readings (e.g., sensor delay and noise/error). The physical models are linked to pseudo components, and the microprocessor is represented as the hierarchy of the pseudo components.

- **Pseudo package:** This pseudo component represents the microprocessor packaging and cooling interface. Temperature and cooling models are linked to the pseudo package, and the temperature is computed at this level based on the power density inputs to the thermal grid.
- **Pseudo partition:** The power is calculated at microarchitecture modules, whereas the temperature is computed at the package level. Therefore, there must be an intermediary component to deliver the module-level power calculation results to the processor package for the thermal simulation. Such matching component is called *pseudo partition*. A partition is physical piece of the silicon die in the processor, which becomes the power source of the thermal RC grid for the temperature modeling. One or more modules may belong to the same partition to control the simulation granularity, depending on the technological, functional, and locational similarity. For instance, the instruction decoder and buffer in the front-end of core pipeline can be grouped into the same partition. Such grouping removes the necessity for modeling tiny modules as power sources and improves the simulation speed. A pseudo package can include arbitrary number of pseudo partitions as well to manage the simulation granularity. The device and technology properties (i.e., voltage, frequency) are assumed to be identical within a partition, and those properties are applied identically to all modules belonging to the partition. The degradation and reliability are computed at this level, based on the device and technology parameters, temperature from thermal models, and switching activity information from the microarchitecture modules.
- **Pseudo Module:** The microarchitecture can be drawn as the linked set of functional components as in Figure 1. Those functional components are called modules, and the pseudo modules are the abstract representation in the

simulation interface. The energy models are linked to pseudo modules, and the energy/power is computed at this level, based on the activity counts collected from architecture simulations.

- **Pseudo Sensor:** All pseudo components in the EI include the data queue structure. The data queue can include multiple types of runtime queues, identified by data name and type. The computation results of the physical models are stored in the data queues of the corresponding pseudo components. A pseudo sensor can attach to any data queues of other pseudo components to monitor the data and emulate the sensor reading. When a new data is inserted into the observed data queue, the pseudo sensor adds the delay and noise to the original data and stores the output in its own data queue which is separate from the monitored data queue. Thus, the original data and read data are separately stored, ensuring the intact data is also kept in the monitored data queue.

Individual models are wrapped into modeling library classes to standardize the interface and connect to appropriate pseudo components. The modeling libraries used in the EI are *energy*, *thermal*, *reliability*, and *sensor libraries*, and each library has own standard functions (e.g., power calculation, temperature computation, failure probability computation), depending on the purpose of models.

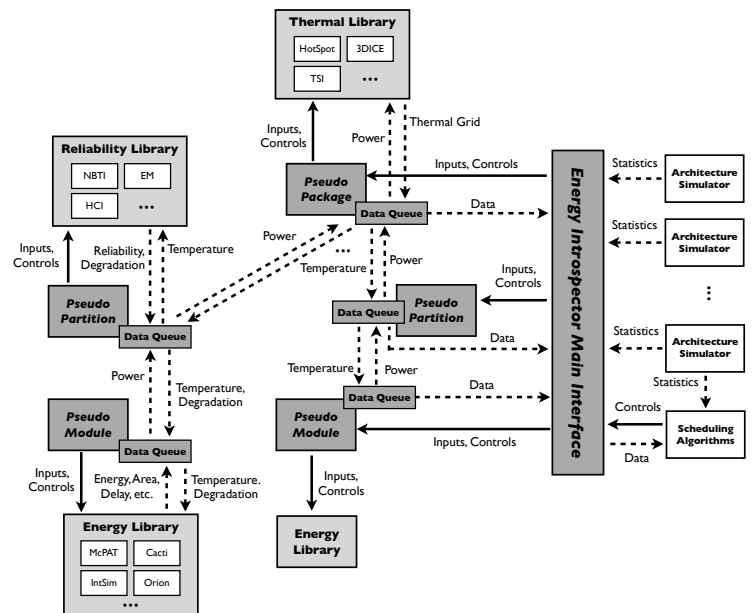


Figure 4: Interactive architecture simulation framework via the Energy Introspector for physical phenomena analyses of the 3D IC.

Figure 4 shows the interactive simulation interface of the EI. The workload characteristics (i.e., switching activity counts) are first collected from the architecture simulations [16]. The collected statistics are passed through the main interface to pseudo modules. The pseudo module use the standard energy library functions to call the specific energy models [17, 18, 19, 20]. Since the pseudo modules have separate links to energy models, the EI allows the

mixed used of energy models by selectively linking the most appropriate models. Calculated power results are collected at pseudo partitions and delivered to the pseudo package. The power numbers are converted to power densities based on the size and location information of pseudo partitions. The power densities are used for the thermal models associated with cooling models to compute the transient change of temperature [2, 21, 22]. The thermal models produce the temperature grid of the silicon layers in the 3D processor package. The temperature result is fed to pseudo partitions to compute the degradation (i.e., threshold voltage shift) and failure probability. [4, 5, 6, 7, 8, 10, 11, 12]. The temperature result from the thermal/cooling models and degradation information from the reliability models are used in the pseudo modules to update the parameters of power calculation at the next sampling time point. Through these processes, the EI accurately captures the interactions among physical phenomena while supporting the runtime architecture simulation. Figure 5 shows the non-uniform failure distribution of an asymmetric 64-core processor layer after executing a set of SPEC2006 benchmarks.




Figure 5: The failure distribution of an asymmetric 64-core processor layer, showing 25.1% of peak-to-peak difference.

3. CONCLUSION

3-dimensional stacking of processing units is a promising solution to continue the growth of high performance computing. The integration of liquid cooling with 3D IC can effectively solve the heat problem, by significantly improving the cooling efficiency compared with air cooling. However, the design space exploration of the 3D IC needs to be carefully conducted under the tradeoffs of performance, energy/power, thermal/cooling effects. And most importantly the reliability must be taken into account, where higher transistor, power, and thermal density per unit volume lead to the significant reduction of the microprocessor lifetime reliability.

4. REFERENCES

- [1] Coskun et al., "Modeling and Dynamic Management of 3D Multicore Systems with Liquid Cooling," *VLSI-SoC*, Oct. 2009.
- [2] Dridhar et al., "3D-ICE: Fast Compact Transient Thermal Modeling for 3D-ICs with Inter-Tier Liquid Cooling," *ICCAD*, Nov. 2010.
- [3] Mizunuma et al., "Thermal Modeling for 3D-ICs with Integrated Microchannel Cooling," *ICCAD*, Nov. 2009.
- [4] Srinivasan et al., "The Case for Lifetime Reliability-Aware Microprocessors," *ISCA*, June 2004.
- [5] Srinivasan et al., "Lifetime Reliability: Toward An Architectural Solution," *Micro*, June 2005.
- [6] Shin et al., "A Framework for Architecture-level Lifetime Reliability Modeling," *DSN*, June 2007.
- [7] Takeda et al., "An Empirical Model for Device Degradation Due to Hot-Carrier Injection," *EDL*, Apr. 1983.
- [8] McPherson, "Reliability Physics and Engineering, Time-to-Failure Modeling," Springer, Aug. 2010, pp. 95-108.
- [9] Leemis, "Reliability: Probabilistic Models and Statistical Methods," 2nd Ed., pp. 127-147.
- [10] Wang et al., "The Impact of NBTI Effect on Combinational Circuit: Modeling, Simulation, and Analysis," *Trans. on VLSI Systems*, Feb. 2010.
- [11] Gupta et al., "GNOMO: Greater-than-Nominal Vdd Operation for BTI Mitigation," *ASP-DAC*, Feb. 2012.
- [12] International Technology Roadmap for Semiconductors, "2011 Update Overview and Table", 2011.
- [13] Raghavan et al., "Computational Sprinting," *HPCA*, Feb. 2012.
- [14] Esmaelizadeh et al., "Dark Silicon and The End of Multicore Scaling," *ISCA*, June 2011.
- [15] Goulding-Hotta et al., "The GreenDroid Mobile Application Processor: An Architecture for Silicon's Dark Future" *Micro*, Dec. 2011.
- [16] Loh et al., "Zesto: A Cycle-Level Simulator for Highly Detailed Microarchitecture Exploration," *ISPASS*, Apr. 2009.
- [17] Li et al., "McPAT: An Integrated Power, Area, and Timing Modeling Framework for Multicore and Manycore Architectures," *Micro*, Dec. 2009.
- [18] Thoziyoor et al., "A Comprehensive Memory Modeling Tool and Its Application to the Design and Analysis of Future Memory Hierarchies," *ISCA*, June 2008.
- [19] Sekar et al., "IntSim: A CAD Tool for Optimization of Multilevel Interconnect Network," *ICCAD*, Nov. 2007.
- [20] Wang et al., "Orion: A Power-Performance Simulator for Interconnection Networks," *Micro*, Dec. 2002.
- [21] Huang et al., "HotSpot: A Compact Thermal Modeling Methodology for Early-Stage VLSI Design," *Trans. on VLSI Systems*, May 2006.
- [22] Cho et al., "Thermal System Identification: A Methodology for Post-Silicon Characterization and Prediction of The Transient Thermal Field in Multicore Chips," *Semi-Therm*, Mar. 2012.