Pathfinder3D: A Framework for Exploring Early Thermal Tradeoffs in 3DIC

Shivam Priyadarshi*, W. Rhett Davis, Paul D. Franzon
Dept. of ECE, Box 7911, North Carolina State University, Raleigh NC 27695

***Abstract*— Three dimensional integration technologies offer significant potential to improve performance, performance per unit power and integration density. However, increased power density and thermal resistances leading to higher on-chip temperature is imposing several implementation challenges and restricting widespread adaptation of this technology. This necessitates the need for CAD flows and tools facilitating early thermal evaluation of possible 3D design choices and thermal management techniques. This paper presents a CAD flow and associated framework called Pathfinder3D, which facilitates physically-aware system-level thermal simulation of 3DICs. Usage of Pathfinder3D is shown using a case study comparing thermal profiles of 2D and three 3D implementations of a quad-core chip multiprocessor.**

Keywords—3DIC, electrothermal simulation, Pathfinding

I. INTRODUCTION

Three-dimensional Integrated Circuits (3DICs) can facilitate higher device density and higher chip-to-chip bandwidths, at lower power levels, than can be achieved with conventional packaging. 3DICs utilize vertical dimension for stacking different ICs. The vertical stacking reduces the interconnect length and routing congestion compared to a conventional 2D implementation and thus reducing interconnect delay and power consumption [1, 2]. Furthermore, 3DICs can broaden the horizon of what s system-on-chip can achieve through heterogeneous integration which means capability to integrate technologies supporting different devices such RF, high performance logic etc. on a single die. The heterogeneous integration can reduce the delay and power consumption and thus allows building energy efficient systems [3, 4]. Moreover, even within one technology, different generations for (example 45nm and 32nm CMOS logic) can be stacked to realize the cost benefits from the better yield of the mature process node [5].

There are four key design challenges associated to 3DICs, which must be addressed for the widespread adaptation of this technology. These challenges are associated to a) heat dissipation, b) power delivery network design, c) floorplanning, and d) design for test. This paper is focused on the thermal challenges of 3DICs. Increased volumetric density in 3DICs leads to large heat-fluxes, and the lower thermal conductivities of the inter-tier and inter-metal dielectrics restrict the heat flow towards the heatsink, making heat removal a challenging task [6, 7]. Moreover, die thinning reduces the amount of silicon and increases the proportion of oxide and molding materials whose thermal conductivities are lower than silicon. These trends result in increased on-chip temperature, which can adversely affect performance (by means of mobility degradation), power (by means of exponential increase in leakage current), reliability (by means of electromigration, time-dependent dielectric breakdown etc.) and cost (by means of increased cooling cost) of packaged 3DICs. Hence, careful thermal design facilitated by modeling and simulation is essential for successfully designing cost-effective high performance 3DICs. Moreover, it is important to address the thermal issues in early design because it provides more opportunities for optimization in low design effort. For example, a study by LSI logic shows that power can be reduced by 20%, 10% and 5% by optimizations at register-transfer level (RTL), gate-level and transistor-level respectively whereas 80% reduction can be achieved at electronic system-level (ESL) [8]. Here "Early" design refers to making technology choices (such as 2.5D vs. 3D, different 3D bonding methods, stacking of different technologies e.g. RF, digital logic), determining an appropriate partitioning (such as homogenous vs. heterogeneous design partitioning, number of stacking layers) amongst the constituent components and determining an approximate floorplan of functions within those components. For studying the thermal tradeoffs among the early designs, physically-aware system/architecture-level CAD flows and tools are required. This paper presents such a CAD flow and a case study showing its application.

The paper is organized as follows. Section II presents the system-level electrothermal simulation flow. Detailed descriptions of various components of flow are presented. Section III presents a case study of a quad-core chip multiprocessor (CMP) system. Section IV concludes this work.

II. SYSTEM-LEVEL ELECTROTHERMAL SIMULATION

The full exploitation of the benefits of 3D integration requires a system-level exploration flow, which can facilitate in finding an optimal 3D design by thermally comparing possible early design choices and thus eliminating thermally bad design early on. This can significantly reduce the development cost and risk. The thermal simulation of an integrated circuit requires some granularity of physical and power details based on which it can be categorized in two groups namely a) fine-grained, and b) coarse-grained simulation. A fine-grained thermal simulation requires accurate physical details such as transistor-level layout, dielectric and interconnects material details (e.g. thermal conductivity, specific heat capacity, density etc.), and power associated to each or group of transistors. Lack of some of these details (e.g. layout, power associated to transistors)

Shivam Priyadarshi is currently with Qualcomm Technologies Inc., Raleigh.

early in the design flow and enormous computational cost of fine-grained simulation prohibit its use in system-level design space exploration. Thus, at system-level, a coarse-grain thermal simulation flow, which allows abstracting the physical details and estimating the power consumption at higher-level of design abstraction, is required. In this section such a system-level electrothermal flow is presented.

A. System-level flow

Fig. 1 [9] shows a system-level 3D design exploration, which contains two parts: the frontend and the backend. In the frontend, the designer creates a high-level system description (using component IP blocks) to evaluate power and performance and provide input to back end for thermal simulation. The framework implementing the flow shown in Fig. 1 is called Pathfinder3D [9].

Transaction-level modeling (TLM) methodology is used to capture the dynamic effects (e.g. transient variations in power profile) resulting from the complex interaction between the system components (i.e. IP blocks) while running an application. Transaction-level modeling is a technique, which separates the computation and communication of a system and hides the details not required at the early phases of the design flow resulting in fast simulation. In a TLM representation, IP blocks contain concurrent processes that execute their behaviors whereas communication is abstracted from cycle-by-cycle operation to an abstract operation called transaction. Communication is implemented as channels, which hide protocols from the IP blocks, and transaction is initiated by calling the interface functions of channels. TLM based simulation can greatly reduce the simulation time compared to Register transfer level (RTL) simulation with acceptable timing accuracy [10] and thus suitable for system-level design space exploration. The authoring of TLM simulations is the most time-consuming part of this flow, and therefore users will most often need to reuse existing TLM frameworks. In this work, Mentor Graphics Vista Framework for SystemC TLM simulation is used [11].

A high-level power model is integrated with TLM framework to calculate the transient power, average power and performance based on the IP configurations. In this paper, a processor-based system is considered thus a statistical modeling approach for predicting the transient and average power of out-of-order superscalar processors is presented.

The backend of the flow starts with a rough floorplan obtained from system-level description using an area model of components. Users must specify technology information of each layer/tier of the stack (the wafer technologies), information about TSV-based stacking of different tiers (a stack technology) and material properties. To speedup the thermal simulation, in Pathfinder3D, layers in the wafer technologies are collapsed into a reduced set of layers, called composite layers. The composite layers, rough floorplan and power information are fed into a thermal simulator to generate the static and dynamic temperature profiles. The system description can be easily changed based on feedback from a cost optimization function taking power, performance and thermal results as the inputs. This flow facilitates fast design space exploration to find an optimal 3D design without doing detailed implementation of all the design choices. In these explorations trends are more important than absolute values.

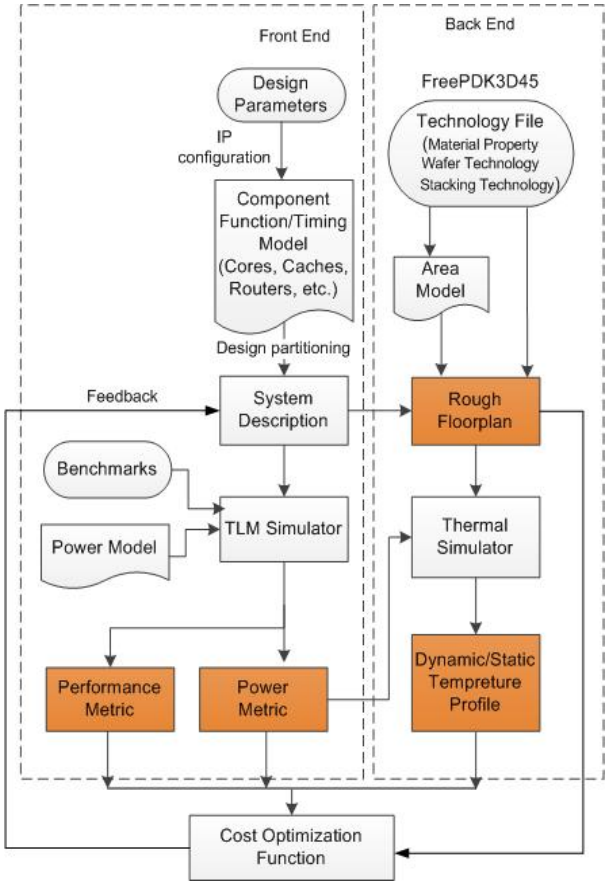

Figure 1: System-level CAD flow for 3D design space exploration.

B. Statistical Power Modeling

Out-of-order (OoO) superscalar processors offer a vast design space consisting of several microarchitectural parameters such as pipeline depth, superscalar width, cache sizes, queue sizes, etc., each with a wide range of potential configurations. It is almost impossible to characterize the whole microarchitecture design space on RTL-level in realistic simulation times. Even exploring this huge design space using C++-level cycle-accurate simulation is infeasible if not impossible. Long simulation times necessitate approaches such as statistical modeling, which reduces the computational cost of design space exploration by significantly reducing the required number of detailed simulations. There can be two types of statistical power model of a processor - a) model predicting average power, b) model predicting transient power. Workloads exhibit a complex dynamic behavior and their runtime characteristics cannot be captured using the average power model [12]. Furthermore, their transient characteristics vary widely for different microarchitecture configurations. Predicting the transient power behavior is similar to predicting

a time series. This requires model construction and training, which can be done in the two steps. As a first step, the transient power traces obtained by running a workload on a small set of sampled microarchitecture configurations should be decomposed into a small set of coefficients capturing the workload power dynamics. Then, in the second step, statistical models should be constructed to predict those coefficients with microarchitectural parameters taken as explanatory variables. In this work, linear, polynomial, and radial-basis function (RBF) based regression techniques are evaluated to construct the statistical model. These models are workload specific, thus, aforementioned two steps need to be repeated for every workload. Once the model is constructed, for any unseen microarchitecture configuration, coefficients can be predicted using the statistical models. Then, the transient power trace can be produced by reconstructing the time series using the predicted coefficients. Once the transient switching power is predicted, leakage power can be added to it and then total power can be fed to Pathfinder3D for dynamic electrothermal simulation.

In this work, eight microarchitecture variables are considered as explanatory variables each taking a range of values. These are superscalar width (fetch and issue width), pipeline depth, ROB (active list), issue queue, load/store queue, L1 data cache, L1 instruction cache, and L2 cache size. This created a design space of 43204 core configurations. Total of 200 core configurations are used for the model construction and training, and 40 core configurations are used for testing (different from 200 used for the model construction). The prediction techniques are evaluated for 9 SPEC CPU2000 workloads. Each workload is fast-forwarded by 1 billion instructions and then cycle-accurate simulation for 300 million instructions is performed. A transient power trace containing 128 samples is generated from the cycle-accurate performance coupled power simulation. The sampling time is taken as 800 us. 32 wavelet coefficients are used to capture the power dynamics of 128 samples. Prediction model for these 32 coefficients are constructed using linear, polynomial and RBF based regression techniques (please refer [13] for more details about each technique). The RBF based regression achieves best prediction accuracy among the three regression techniques explored in this work. Root mean square error (RMSE) in the transient switching power prediction using RBF based regression is plotted as boxplot in Fig. 2. The lower and upper end of the central box represents the first and third quartile respectively. The line within the central box represents median. Whiskers represent the extreme data values that are not an outlier. The distribution of 40 root mean square errors expressed in percentage obtained by predicting the transient power trace of 40 cores elected for testing the accuracy of regression. The line with dot shape markers shows mean of percentage RMSE across all the 40 test cores. Average of mean RMSE across all the 9 workloads (i.e. average of 9 dot shape markers) come out to be 7.6%. Fig. 3 shows the error in predicting the average power. It shows that for most of the workloads, RBF based regression approach predicts average power within 5% error bound.

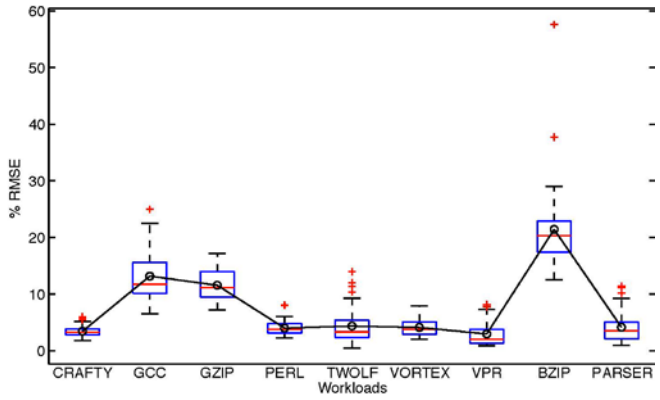

Figure 2: Boxplot representing the percentage root mean square error in transient switching power prediction using RBF based regression for different workloads.

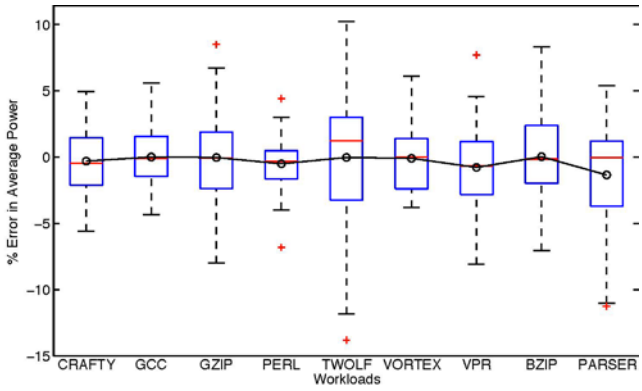

Figure 3: Boxplot representing the percentage error in average power prediction using RBF based regression for different workloads.

C. Electrothermal Simulation

As shown in Fig. 1, a rough floorplan and power profiles (average and transient) of floorplan blocks are required for the thermal simulations. Pathfinder3D allows a user to specify a textual description of their floorplan. The user can define the dimensions of basic building blocks of each tier in the 3D stack as macrocells. Each macrocell can have multiple sockets representing ports of a block. A user can then replicate these macrocells as different instances at various locations in their corresponding tiers by specifying their co-ordinates and the connections among instances. The instance connection information can be imported from the system description defined in the front-end. A simple routing tool then can be applied to estimate the global wire length. Furthermore, number of repeaters can be estimated using a minimum delay insertion algorithm. This allows estimating the interconnect power of the design.

For the steady-state thermal simulation, users can specify the average power of each floorplan instance in a separate power stimulus file at the beginning of the simulation. Pathfinder3D reads the power stimulus file and

assigns the power of each instance to the active layer of that instance. Currently, power is uniformly distributed across the area of a macrocell. For the dynamic electrothermal simulation, the TLM and thermal simulators are integrated and synchronized in a relaxation loop within the Pathfinder3D framework. A lock-step synchronization algorithm managed by a global synchronizer is used for the electrical-thermal co-simulation. After every *Tsample* transactions, the TLM simulator generates a transient power trace in the stimulus file. At this point, the global synchronizer suspends the SystemC TLM simulation and starts the thermal simulation using the power trace generated in the TLM simulation. After the transient thermal simulation for the time duration of *Tsample* transactions, the global synchronizer stops the thermal simulation. At this point, thermal simulator feeds back the current temperature of the component IP blocks to the TLM simulator where the leakage power is updated based on the temperature. Now, the global synchronizer resumes the SystemC TLM simulation. This synchronization mechanism is illustrated in Fig. 4. The synchronizer uses SystemC *sc_start(<time)* function to suspend and resume the TLM simulation. When called the first time, this function will start the SystemC scheduler, which will run up to the simulation time passed as an argument. When called on the second and subsequent occasions, this function will resume the scheduler from the time it had reached at the end of the previous call to *sc_start()*. The scheduler will run for the time passed as an argument, relative to current simulation time. In this way, simulation time maintained by the SystemC kernel does not increase during the thermal simulation. Due to the fact that thermal transient effects are typically much slower than electrical effects, *Tsample* is typically kept much larger than the amount of time needed for *Isample* instructions (power is calculated on the granularity of *Isample* instructions). Note that once the thermal simulator receives a start signal from the synchronizer, it becomes independent of the TLM simulator and it can perform the transient analysis on its own timing granularity (i.e., time step size) based on the speed verses accuracy tradeoffs.

The Pathfinder3D toolset consists of a physical thermal extractor, WireX [14]. WireX reads the layout generated from the rough floorplan and creates a linear resistive and capacitive thermal netlist. It uses the thermal conductivity, specific heat capacity and other material properties of the composite model to generate this netlist. WireX meshes the layout and discretizes it in cuboids. Each cuboid is modeled using a thermal resistor from the center of cuboid to each face with a thermal capacitor from center terminal to the thermal ground. The user is able to control the fidelity of the mesh. Usually for fast thermal simulation in pathfinding phase the resolution of mesh is kept low. For steady state thermal analysis, the extractor generates a thermal modified nodal admittance matrix (**Y**) and power vector (**J**). Then, **Yv** = **J** is solved by a linear sparse matrix solver to get the temperature vector (**v**). The WireX tool currently supports only one style of boundary condition, which is a perfect heatsink on one face of the stack and adiabatic on all other faces. The perfect heatsink is assumed to be connected to the first tier specified in the stack technology. All temperatures will be reported as rise above the heatsink. For absolute temperatures, the temperature calculation must take the package and physical heatsink in account. The current model is accurate assuming the substrate of first tier is a perfect heat spreader which is not always an accurate assumption. However, users have flexibility to model portions of the package and heatsink as composite layers in the Pathfinder3D technology file.

The transient thermal simulation requires a transient solver. Currently Pathfinder3D does not include any transient solver. However, it can generate a netlist compatible with HSPICE or the open source circuit simulator fREEDA [15]. Simulation speed can be significantly increased using the parallel transient simulation technique for multiphysics circuits [16] available in fREEDA.

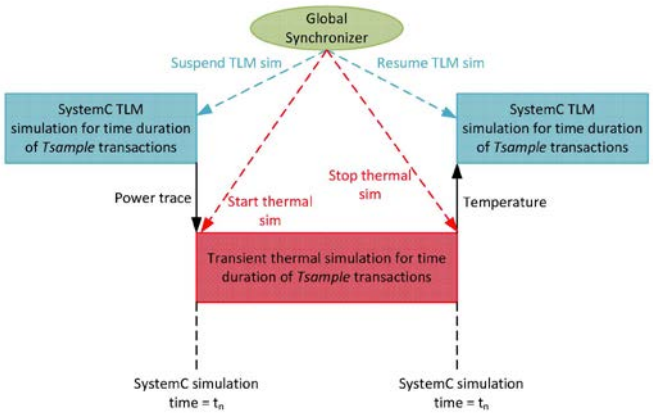

Figure 4: Lock-step synchronization mechanism for the electrothermal simulation.

III. CASE-STUDY

To demonstrate the applicability and usefulness of flow presented here, 2D and three 3D implementations of a quad-core CMP are considered. Fig. 5 shows the 2D floorplan (called 2D FLP) of the quad-core CMP. Each core in the system is of dimension 1.52**x**2 mm and consumes 1.02 W of power on average and represents a four wide out-of-order superscalar processor. Total L2 cache size is 2 MB, which is distributed in four banks as shown in Fig. 5. Dimension and average power of each L2 bank are 2.36**x**2 mm and 100 mW respectively. A crossbar style router for establishing the connection between the cores and L2 cache banks is used. The dimension and power of router are 0.5**x**0.8 mm and 20 mW respectively.

Three 2-tier 3D implementations of the quad-core CMP namely a) 3D FLP1, b) 3D FLP2, and 3D FLP3 are considered here. In 3D FLP1, at each tier there are two cores and two L2 cache banks. In this floorplan, core on Tier B is stacked over core on Tier A where Tier A is near the heatsink. Similarly L2 cache bank on Tier B is stacked over L2 cache bank on Tier A. The router is located on Tier A. TSVs are

used for connecting cores and L2 cache banks on Tier B with router on Tier A.

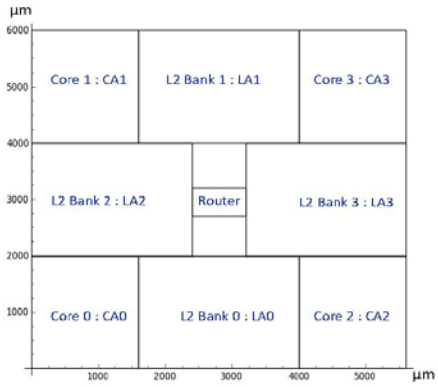

Figure 5: Floorplan of quad-core 2D CMP considered in this case study.

3D FLP2 is similar to 3D FLP1 except in this case, cores on Tier B are stacked over L2 cache banks on Tier A, and L2 cache banks on Tier B are stacked over cores on Tier A. In 3D FLP3, all the four cores and router are located on Tier A. Tier B has all the four L2 cache banks. This represents SRAM over logic case. Table 1 presents a summary of results obtained by comparing 2D and 3D implementations in terms of maximum temperature rise and maximum spatial temperature gradient across a grid block. It shows, 3D FLP3 (SRAM over logic case), exhibits lowest temperature rise and lowest spatial temperature gradient. Fig. 6 shows the static thermal profile of 3D FLP3. In the thermal simulations presented here, a heat sink based cooling solution with forced air-convection is considered. A single lumped thermal convection resistance of 0.8 K/W is used. This case study shows that using the flow presented in Section II, comparisons like shown in Table 1 can be quickly performed at system-level without any detailed implementation guiding system architects in eliminating thermally bad design choices early on.

TABLE 1: THERMAL COMPARISON OF 2D AND 3D FLOORPLANS

Metric	2D FLP	3D FLP1	3D FLP2	3D FLP3
Max. Rise in temp.	26.2 K	39.8 K	30.5 K	28.25 K
Max. Spatial temp. Gradient	4.02 K	4.51 K	2.19 K	1.26 K

Fig. 7 shows the dynamic thermal profile of 2D and 3D floorplans. A transient power trace is required for the dynamic simulation, which is obtained by executing *gcc* workload from SPEC CPU2000 on the cores on Tier A and *bzip* workload from SPEC CPU2000 on the cores on Tier B on a TLM simulator. Cores in this case study are twelve deep four wide OoO superscalar processor, with 32 KB data and instruction cache, 96 entries ROB, 32 entries issue queue, and 160 entries physical register file running at 1.5 GHz. The dynamic power trace of other processor microarchitecture configurations can be generated using the statistical power modeling approach presented in Section II. Fig. 7 also illustrates that floorplan 3D FLP3 is thermally better than other 3D floorplans, 3D FLP1 and 3D FLP2. Furthermore, 3D FLP1 is significantly hotter than other floorplans. Temperature dependence of leakage power is not considered in the dynamic simulations presented here. Considering leakage-temperature positive feedback can produce different trend than shown in Fig. 7.

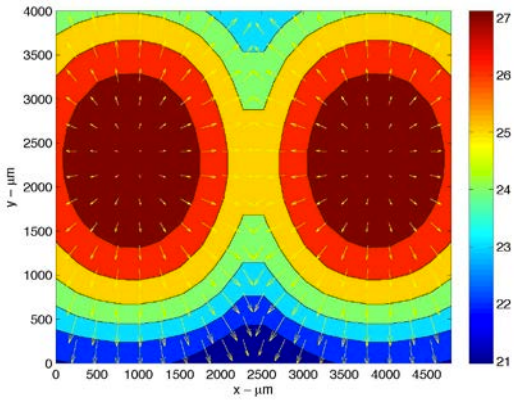

Figure 6: Spatial thermal profile corresponding to Tier B of 3D FLP3.

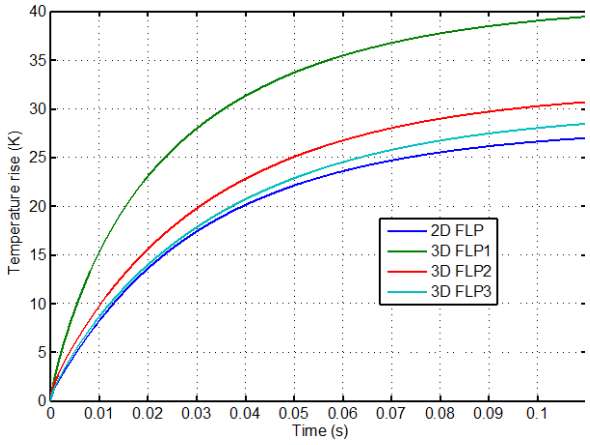

Figure 7: Dynamic thermal profile of 2D and 3D floorplans of quad-core CMP.

IV. CONCLUSION

Designing an optimal 3DIC requires CAD tools and flows to explore the tradeoff between various available design choices early in the design cycle. This paper presents a flow for the comparison of static and dynamic thermal profile of possible design choices without the detailed implementation. Thermal profiles of 2D and 3D quad-core chip multiprocessors are compared. One of the key challenges of the system/architecture level thermal simulation is unavailability of power models for all possible design choices early on. To

address this issue, a statistical power modeling technique for out-of-order superscalar processor is presented. Furthermore, system-level synchronization mechanism between TLM and thermal simulation is discussed in detail.

REFERENCES

[1] Davis, W. R., Wilson, J., Mick, S., Xu, J., Hua, H., Mineo, C., Sule, A., Steer, M. B., and Franzon, P. D., 'Demystifying 3D ICs: The pros and cons of going vertical', IEEE Design & Test of Computers, Nov 2005, 22, (6), pp. 498-510.

[2] Banerjee, K., Souri, S., Kapur, P., and Saraswat, K., '3-D ICs: a novel chip design for improving deep-submicrometer interconnect performance and systems-on-chip integration', Proc. Of the IEEE, May 2001, 89, (5), pp. 602-633.

[3] Bernstein, K., Andry, P., Cann, J., Emma, P, Greenberg, D., Haensch, W., Ignatowski, M., Koester, S., Magerlein , J., Puri, R., and Young, A., 'Interconnects in the Third Dimension: Design Challenges for 3D ICs', Proc. IEEE/ACM Design Automation Conference, June 2007, pp. 562-567.

[4] Franzon, P. D., Rotenberg, E., Tuck, J., Davis, W.R., Huiyang Zhou, Schabel, J., Zhang, Z., Park, J., Dwiel, B., Forbes, E., Joonmoo Huh, Priyadarshi, S., Lipa, S., Thorolfsson, T., 'Applications and design styles for 3DIC', Proc. *IEEE International Electron Devices Meeting (IEDM)*, pp.29.4.1-29.4.4, Dec. 2013.

[5] Priyadarshi, S., Choudhary, N., Dwiel, B., Upreti, A., Rotenberg, E., Davis, W. R., and Franzon, P. D., 'Hetero^2 3D integration: A scheme for optimizing efficiency/cost of chip multiprocessor', Proc. International Symposium on Quality Electronic Design, March 2013.

[6] Cong, J., Wei, J., Zhang, Y., 'A thermal-driven floorplanning algorithm for 3D ICs', Proc. IEEE/ACM International Conference on Computer Aided Design, pp. 306-313, 7-11 Nov. 2004.

[7] Priyadarshi, S., Harris T., Melamed, S., Otero, C., Kriplani, N., Christoffersen, C., Manohar, R., Dooley, S., Davis, W., Franzon, P., and Steer, M., 'Dynamic electrothermal simulation of three-dimensional integrated circuits using standard cell macro-models', IET Circuits, Devices Systems, vol. 6, no. 1, pp. 35-44, January 2012.

[8] Higher abstraction saves more power. [Online]. Available: http://www.iqmagazineonline.com/current/pdf/Pg30-33_IQ_32-Higher_Abstraction_Saves_More_Power.pdf

[9] Priyadarshi, S., Hu, J., Choi, W. H., Melamed, S., Chen, X., Davis, W. R., and Franzon, P. D., 'Pathfinder 3D: A Flow for System Level Design Space Exploration', Proc. IEEE International 3D System Integration Conference, Feb 2012, pp. 1-8.

[10] G. Beltrame, D. Sciuto, C. Silvano, 'Multi-Accuracy Power and Performance Transaction-Level Modeling', IEEE Transactions on Computer-Aided Design of Integrated Circuits and System, vol.26, no.10, pp.1830-1842, Oct. 2007.

[11] Mentor graphics vista flow. [Online]. Available: http://www.mentor.com/esl/vista/flow

[12] Cho, C., Zhang, W., and Li, T., 'Informed microarchitecture design space exploration using workload dynamics', in 40th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), 2007, pp. 274-285.

[13] Priyadarshi, S., 'System and Gate-Level Dynamic Electrothermal Simulation of Three Dimensional Integrated Circuits', PhD dissertation, North Carolina State University, 2013.

[14] Melamed, S., Thorolfsson, T., Harris, T. R., Priyadarshi, S., Franzon, P. D., Steer, M. B., Davis, W. R., 'Junction-Level Thermal Analysis of 3-D Integrated Circuits Using High Definition Power Blurring', IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, May 2012, 31, (5), pp. 676-689.

[15] freeda : An open-source multiphysics simulator. [Online]. Available: http://www.freeda.org/

[16] Priyadarshi, S., Saunders, C., Kriplani, N., Demircioglu, H., Davis, W., Franzon, P., and Steer M., 'Parallel transient simulation of multiphysics circuits using delay-based partitioning', IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 31, no. 10, pp. 1522–1535, Oct 2012.