Early Chip Planning Cockpit

Jeonghee Shin, John A Darringer, Guojie Luo†, Alan J Weger and Charles L Johnson
IBM T. J. Watson Research
Yorktown Heights, NY 10598
jeonshin@us.ibm.com

***Abstract*—The design of high-performance servers has always been a challenging art. Now, server designers are being asked to explore a much larger design space as they consider multicore heterogeneous architecture and the limits of advancing silicon technology. Bringing automation to the early stages of design can enable more rapid and accurate trade-off analysis. In this paper, we introduce an Early Chip Planner which allows designers to rapidly analyze microarchitecture, physical and package design trade-offs for 2D and 3D VLSI chips and generates an attributed netlist to be carried on to the implementation stage. We also describe its use in planning a 3D special-purpose server processor.**

Keywords-system level design automation; early chip planning

I. INTRODUCTION

In a highly competitive market, server designers are continually pushing the limits of design and technology. As a result, power and voltage drop are always at the limits. Now with the emergence of 3D [1], which enables stacked chips to be interconnected with through-silicon vias (TSVs), there are more options and issues to deal with. At the same time, server architecture has been exploring multiple cores for high throughput and heterogeneity for single-thread performance improvement [2]. The task of designing these multicore heterogeneous systems in a 3D technology presents the designers with an explosion of trade-offs to consider. As a result, the automation of early chip planning has become more vital than ever for efficient design space exploration [3]. Our work is motivated by this need, especially in the server domain.

Today, server system designers often use a combination of Excel spreadsheets and PowerPoint diagrams to consider alternatives. They are usually limited in the number of choices they can effectively consider, and the quality of the results depends on their experience. In addition, the new chip often needs to be designed for more than one system configuration (e.g., high-end and low-end), requiring different design constraints of frequency, power, temperature or even lifetime. Moreover, time-to-market pressure forces a new chip design to be conducted in parallel with technology development and the final phases of the preceding processor's design. The goal of this work is to automate the designer's manual efforts at the early stages and provide an early chip planning environment where they can rapidly analyze microarchitecture, physical and package design trade-offs, while keeping up with evolving technology and design data over time.

The rest of the paper is organized as follows: Section 2 describes the early chip planning environment. Section 3 continues the planning discussion for 3D chips. Section 4 discusses related work and Section 5 concludes the paper with future work.

II. EARLY CHIP PLANNER

Figure 1 provides an overview of the Early Chip Planner (ECP). It has a "front-end" for designer inputs and outputs, and a "back-end" which performs the analysis. Currently, the front-end interface is Microsoft Excel spreadsheets, where the designer can specify architecture configurations, technology scaling rules, and package properties for thermal and power distribution analysis. Built-in Visual Basic macros [4] take the designer inputs, set up a remote connection to the back-end server, initiate the execution of the tools, and wait until the execution completes and transfers the results back to the spreadsheets. Once the results are received, they are displayed in a pop-up window as text and plots. This process requires only a button click in the spreadsheets. The similar interface can be implemented with other applications such as web-based front-end.

The main functions of ECP run on the back-end server as illustrated in Figure 1 and described below.

Reference Design Library: In many cases, a new server chip design exploits the preceding design. Thus, early chip planning begins by establishing a library of these "reference" designs, which includes physical properties such as size, location and pins, and power estimates such as data, latch, array active and leakage power. The level of detail can range from macro-level to chip-level. ECP extracts data from multiple databases in different formats and converts them to an XML-based internal database called PAD (Processor Architecture Description). The automated conversion process keeps up with ongoing reference design changes, which often take place while the new processor is being considered.

Physical Planning: Based on the designer's architecture specification and technology scaling rules, Physical Planner assembles and places the proper units, which are extracted from the reference design library and scaled with respect to target technology and microarchitecture change. At the beginning of design, technology scaling rules are generally available as simple scaling factors for area and power. As the target technology becomes better defined, the scaling rules are refined, considering physical design properties (e.g., device types, size, latch count, frequency and voltage) and workloads (e.g., utilization, switching and clock gating factors).

Design Trade-off Analysis: One of the biggest challenges in designing high-performance server processors is high power density, which creates thermal hotspots, high IR drop and

† Guojie Luo is a PhD student at UCLA. The work was performed while he was an intern at the IBM research center.

978-3-9810801-7-9/DATE11/©2011 EDAA

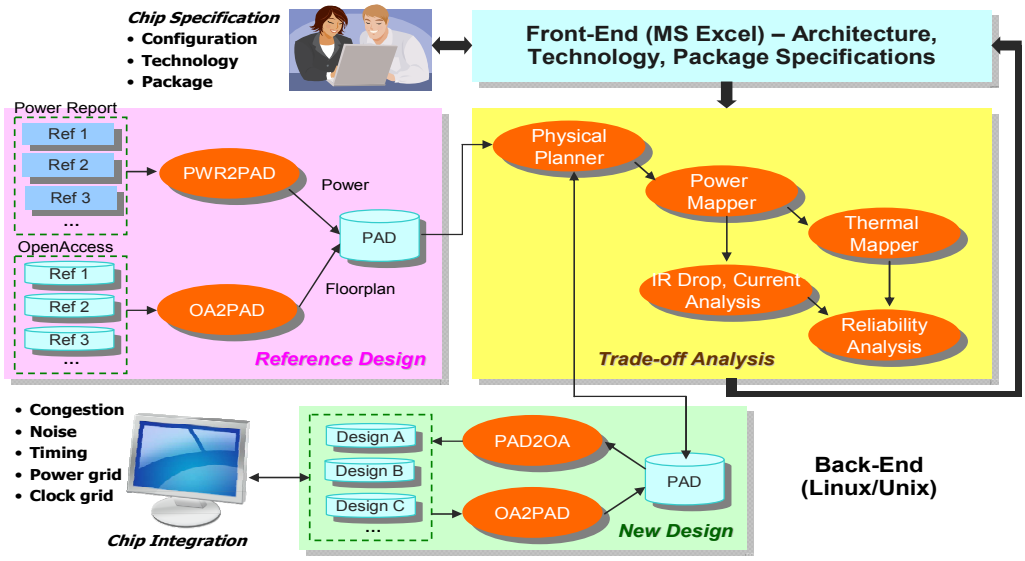

Figure 1. Early Chip Planner (ECP) Environment.

excessive C4 current demand [5,6]. As chip power density tends to increase with technology scaling and 3D technology, these thermal and electrical challenges make early analysis and optimization more critical. In ECP, local power density is analyzed by the Power Mapper at any specified grid size or macro boundaries, using the combined physical design and power data from Physical Planner. For thermal analysis, the Thermal Mapper takes the power map, as well as the package geometry and material properties specified in the front-end spreadsheet, and generates a thermal profile for each chip. The Thermal Mapper incorporates ChipJoule3D [7] to solve thermal networks and the solution is saved in a matrix for each package type. Then, the matrix is used for rapid thermal analysis of alternative chip designs to avoid long simulations [8]. The thermal profile is also used to update leakage power and check thermal runaway. ECP also analyzes IR drop and C4 current for limits and electrical integrity. Further, the generated current distribution, along with the thermal profile is used for chip lifetime estimation with respect to electromigration [6].

Package Planning: Chip packaging is typically planned well ahead of the chip physical implementation to meet a manufacturing schedule. At that time, design inputs for package evaluation such as local power demand is generally unavailable in the manual chip planning process. Thus, a thermal package often assumes worst-case chip power density, which can result in over-design for cooling. Similarly, C4 allocation and placement for electrical package design is often done with average power density. If later analysis shows that any of the C4s exceed the current limit due to higher local power demand, there can be a very expensive design modification required and a delay in product shipment. With ECP, chip power density estimation is available very early in the schedule, enabling architecture, physical and package design to be efficiently evaluated and planned together.

Netlist Generation: A netlist of the new design is created in some detail with data, power and clock pins of the composed units, based on pin information extracted and adjusted from reference design and the physical planning results stored in PAD. Currently, a basic abstract netlist is produced to represent inter-unit interconnection essential for early floorplanning and wire congestion analysis. Power and clock pins are used for power and clock grid planning. For 3D chips, this early netlist generation also provides an estimate of TSV demand along with an initial placement. When the planning process is complete, the generated netlist is passed on to the implementation stage.

III. 3D Chip Planning

In this section, we describe early 3D chip planning with focus of microarchitecture partitioning, TSV planning and chip modularity. The discussion is based on an example of planning a 3D special-purpose server processor with the early chip planning environment introduced in the previous section.

Reference Design: The example 3D processor discussed in this section is based on IBM Wire Speed Power™ Processor, a high-performance network processor [9]. The reference design library includes the physical design and power estimates of its cores, accelerators and I/O units. Figure 2 shows the top-level floorplan and power map of the reference design (fully functioning mode with 2.3GHz and 1V).

3D Physical Planning: In addition to scaling and floorplanning, 3D chip planning requires partitioning of the microarchitecture under consideration across multiple dies. Physical Planner takes the partitioning options specified by the designer and, for each option, does scaling and placement of proper units extracted from the reference library. Currently, Physical Planner does not allow finer-grain partitioning at the unit or macro-level. In planning the example 3D processor, many partitioning and stacking options are considered. Figure 3 illustrates a few examples. Figure 2 shows more details of a 4-high stack with one accelerator die (top), two compute dies (middle) and one IO die (bottom). The four dies are connected with PowerBus [9] at the center of the dies, creating signal TSV farms. Physical Planner effectively estimates TSV counts and area based on the netlist and TSV pitch, and decides the

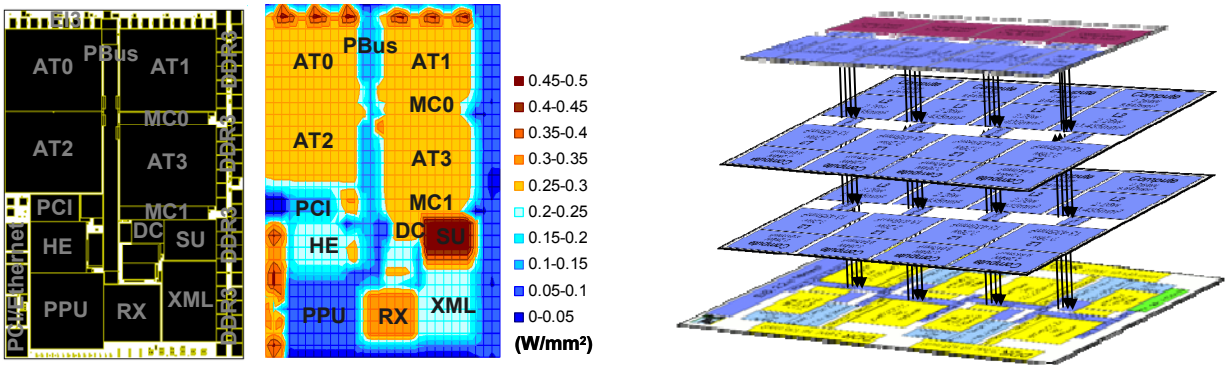

Figure 2. Reference design floorplan (left, 16.7x24.5mm^2), power map (middle, total power: 65W) [9] and 3D baseline design (right, 4 high stack) [10,11].

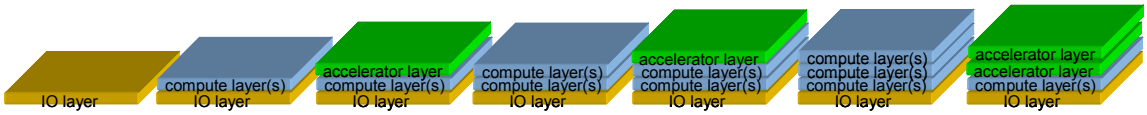

Figure 3. 3D stacking examples [10,11]. One or more dies with a different mix of units can be stacked in one or taller high to reduce the 3D chip height.

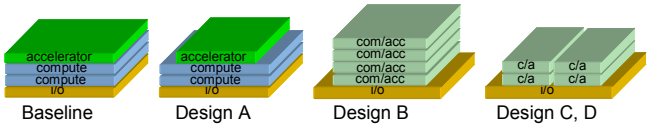

Figure 4. 3D configurations for comparison in Figure 5.

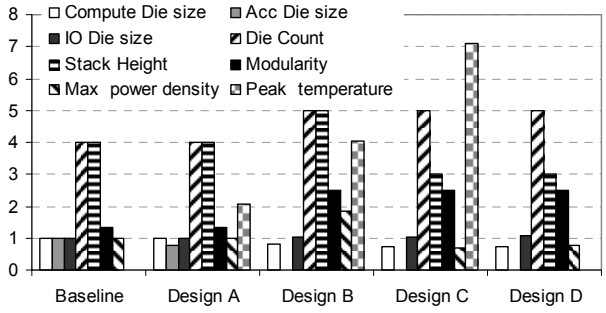

Figure 5. Trade-off analysis results. Die size and maximum power density aggregated for the 3D chip are presented with respect to the baseline design. Peak temperature indicates temperature increase from the baseline design. Modularity is presented as total die count divided by unique die count.

location of the signal TSVs as a part of floorplanning.

Design and Package Trade-off Analysis: Managing power density in the design of server chips has always been a challenge, and 3D technology compounds the problem. Tall chip stacks present both power delivery and cooling issues, thus shorter chip stacks are generally preferred. However, shorter stack configurations can reduce modularity or increase package design complexity. Figure 4 shows 5 of the configurations considered in this section. Figure 5 compares the properties of the alternative configurations.

The baseline design has 3 unique dies stacked 4-high. The compute and accelerator dies are the same size and slightly smaller that the I/O die, which is on the bottom close to the package. The equal upper die size simplified packing, but the accelerator die size can be reduced. Design A has the same partitioning as the baseline, but the accelerator die size is reduced by 22%. This reduces die cost, but results in about 2 degree temperature increase and more complex packaging.

To increase modularity, the system microarchitecture is repartitioned to a mix of cores and accelerators on the same die. This allows four copies of the compute/accelerator die to accommodate the same number of cores and accelerators. As a result, unique die count decreases to two (compute/accelerator and I/O). But, total die count increases to five. Design B achieves the higher modularity by integrating the dies in a 5-high stack, but this leads to higher power density and peak temperature.

To mitigate peak power density, Design C uses two stacks of four compute/accelerator dies and puts them on the I/O die, thus becoming a 3-high chip. However, the thermal analysis of Design C shows a significant increase in peak temperature because the gap between the two stacks is empty in the typical package type and any heat source under the gap area causes hotspots. Design D reduces the peak temperature by removing heat source under the gap area during floorplanning, similar to the baseline design's temperature. Another way to mitigate such thermal hotspots is to specify a specific package that is capable of cooling this gap by filling it with heat conductors.

IV. Related Work

Cadence's Chip Planning System [12] also attempts to satisfy the need for early design automation. It helps designers choose a set of Intellect Property (IP) blocks while considering die size, power and performance for multiple designs. InCyte Chip Estimator [13] allows the exploration of different architectural options and IP selection. However, these tools are aimed at the SoC domain and not intended for 3D technology. In addition, they do not support thermal, electrical or reliability analysis, which is critical for server processors, or consider package properties.

Other related work involves physical planning for system-level design. In particular, studies of 3D architecture

exploration are performed in [14,15], using automatic 3D floorplanning to exploit vertical interconnection. The studies reported the performance and power benefits of 3D technology. While this work can be useful for early chip planning, it is limited to floorplanning and does not provide a comprehensive chip planning environment, as the Early Chip Planner does. The proposed methodologies can be adopted by the Physical Planner for 3D floorplanning.

V. Conclusion and Future Work

The paper describes an Early Chip Planner which improves the productivity of 2D or 3D chip design at the early design phase and provides an initial chip layout to the chip integration environment. The Early Chip Planner has been used in several server design projects and enabled designers to explore many more alternatives with increased accuracy, compared to their traditional manual methods. As discussed in this paper, the close coupling of microarchitecture, physical and package design has allowed new alterative server designs to be evaluated in hours, instead of days or weeks, which would be required in some traditional manual processes. The automatic scaling and floorplanning, based on detailed physical implementation data derived from reference designs, has exposed more subtle issues, which might not have surfaced until much later in chip design. The generation of an initial top-level netlist for the designs in the Open Access format [16] enables the design decisions and constraints to be passed directly to tools in the implementation phase of design.

More work is underway to expand the scope of analysis to include additional important physical properties and to increase the level of optimization in producing alternative designs. Some of the main thrusts for future research are discussed below.

- System and chip performance are currently estimated with a different set of tools. We plan to develop an abstract model derived from the detailed but slow performance models and integrate it into the Early Chip Planner for rapid performance estimates.
- While the Physical Planner currently supports micro-architecture-level partitioning, we plan to enable finer-grain partitioning at the unit or macro-level for future 3D technologies with smaller TSV dimension and pitch. For finer-grain partitioning, floorplanning will need more detailed physical design information, timing and TSV properties, as shown in prior work [15].
- The computation-intensive thermal and electrical analysis generally limits the turn-around time of trade-off analysis and the number of design alternatives that can be evaluated for early chip planning. To address this issue, we are developing faster methods, based on the fact that for many package models a significant part of the computation can be saved and reused without sacrificing accuracy. We are also considering the extension of the analysis capability to include dynamic thermal and voltage drop analysis.
- Using the thermal and electrical profile provided by the Early Chip Planner, chip lifetime can be estimated. Currently, the Early Chip Planner has the capability of estimating chip lifetime due to C4 electromigration failure as mentioned in the paper. We are integrating analytical lifetime reliability models for key failure mechanisms and design structures, proposed in a previous work [17,18].
- Today's Early Chip Planner exploits significant user guidance in generating configurations to be evaluated. Ultimately, we envision a more automated process that generates optimized configurations satisfying the designer's specification. This will require more work on effective configuration generation methods and the use of a constraint solver with area, performance, power, temperature and reliability as its objectives. Further, we need to continue working on fast yet accurate performance, power, thermal and reliability models.

Acknowledgment

We would like to thank our colleagues at IBM who have been working with us and provided invaluable insights; special thanks to Pradip Bose, Alper Buyuktosunoglu and Eren Kursun. We would also like to thank Gero Dittmann who built the initial Early Chip Planner infrastructure.

References

[1] "3D Chip Technology," IBM J of R&D, Vol. 52, No. 6, 2008.

[2] "Hybrid Computing Systems," IBM J of R&D, Vol. 53, No. 5, 2009.

[3] J. A. Darringer, "Multi-Core Design Automation Challenges," DAC, pp. 760-764, June 2007.

[4] Microsoft Excel Language Reference, http://msdn.microsoft.com.

[5] S. Bobba et al., "IC Power Distribution Challenges," ICCAD, pp. 643-650, November 2001.

[6] J. W. Nah et al., "Electromigration in Pb-Free Flip Chip Solder Joints on Flexible Substrates," J of Applied Physics, Vol. 99, No. 2, pp. 023520, 2006.

[7] IBM Electromagnetic Field Solver Suite of Tools, http://www.alphaworks.ibm.com/tech/eip.

[8] H. F. Hamann et al., "Spatially-resolved imaging of microprocessor power (SIMP): hotspots in microprocessors," ITHERM, pp. 121-125, 2006.

[9] C. Johnson et al., "A Wire-Speed Power™ Processor: 2.3GHz 45nm SOI with 16 Cores and 64 Threads," ISSCC, February 2010.

[10] S. Daijavad, "Application and Network Optimized Next Generation System Research," GRC ETAB Summer Study, June 2009.

[11] D. P. LaPotin et al., "Application and network-optimized computing systems," IBM J of R&D, Vol. 54, No. 1, February 2010.

[12] Cadence, Chip Planning System, http://www.cadence.com.

[13] Cadence, InCyte Chip Estimator, http://www.chipestimate.com.

[14] Y. Xie et al., "Design Space Exploration for 3D Architecture," J of Emerging Technologies for Computer Systems, Vol. 2. No. 2, pp.65-103, April 2006.

[15] Y. Liu et al., "Fine Grain 3D Integration for Microarchitecture Design Through Cube Packing Exploration," ICCD, pp. 259-266, October 2007.

[16] The Open Access Standard, http://www.si2.org.

[17] J. Shin et al., "A Framework for Architecture-Level Lifetime Reliability Modeling," DSN, pp. 534-543, June 2007.

[18] J. Shin et al., "A Proactive Wearout Recovery Approach of Exploiting Microarchitectural Redundancy to Extend Cache SRAM Lifetime," ISCA, pp. 353-362, June 2008.