DESIGN AUTOMATION FOR SELF-ADJUSTING ARCHITECTURES

Jieyi Long (jieyi-long@northwestern.edu)
Northwestern University, Evanston, IL 60208, USA
PhD Advisor: Prof. Seda Ogrenci Memik
Estimated Graduation Date: June, 2011
This work has NOT been published at ASPDAC or DATE PhD forum

I. ABSTRACT

Steady miniaturization and large scale integration place new implications on VLSI system design flow. Firstly, as technology advances, power density of modern VLSI systems keeps increasing. Power is converted into heat, causing thermal variation in both temporal and spatial terms. Since temperature has a direct impact on the delay of the CMOS gates, the timing correctness of the system cannot be guaranteed at design time unless designed conservatively. Secondly, as the feature size of the transistors shrinks, process variation has a larger and larger influence on the timing of the circuits. Although various statistical analysis and optimization techniques have been proposed, they are essentially static techniques which are not tailored for individual chips. To reach the performance limit of each individual chip, a better methodology is to develop self-adjusting architectures which can constantly detect the actual operating condition and instantaneously adapt to environmental changes.

In this proposal, we present our work on the design automation for self-adjusting architectures. Essentially, two approaches, namely, global and local feedback loop based self-adjusting mechanisms are considered. The global feedback loop based mechanisms involve a set of sensors, a set of adjustable elements, and one/multiple central control unit(s) that cooperate together to achieve self-adaptation. In contrast, in local feedback loop based mechanisms, a set of specially designed elements can sense the environmental changes and adjust themselves to achieve resiliency for the entire system to the changes. In summary, the global feedback loop based mechanisms perform the self-adjustments at a global scale, while local feedback loop based mechanisms guarantee the global correctness by performing only local self-adjustments. On the global feedback loop track, we have designed a systematic framework for thermal sensor placement in chip multiprocessors, and SAP, a Self-Adjusting Pipeline architecture. On the other track, we have developed SACTA, a Self-Adjusting Clock Tree Architecture.

The complexity of the auxiliary components that support the self-adjusting feature increases as the complexity of the entire design grows. Therefore, effective automated design frameworks for the self-adjusting architectures are highly desired. We have developed several automated design frameworks for our proposed self-adjusting architectures. Experimental results illustrate the effectiveness of our design frameworks. For instance, our sensor placement techniques are able to cover 98.7% of all the thermal emergencies using only 16 sensors per core; applying SAP, the self-adjusting pipeline architecture to a DEC Alpha-like microprocessor pipeline results in 9.5% improvement in batch performance; and SACTA, our proposed self-adjusting clock tree architecture, is able to prevent thermal induced timing violations within a significantly larger range of operating temperatures (enhancing the violation-free range by 15°C on average).

In the following, we provide more details for each of our works.

A. Global Feedback Loop Based Self-Adjusting Mechanism:

Global feedback loop based self-adjusting mechanism involves 1) a set of sensors (can be thermal sensor, delay sensor, leakage sensor, etc.), which are in charge of monitoring the operating condition of the entire system; 2) a set of adjustable circuit elements, which can be adjusted to ensure system resilience to operation condition variations; and 3) one or multiple central control units, which configure the adjustable elements according to the sensor readings. In the following, we first discuss our work on the sensor system design. Then, we propose a Self-Adjusting Pipeline (SAP) architecture, which implements the idea of global feedback loop based self-adjusting mechanism at circuit level.

Accurate Operating Condition Monitoring Scheme

The sensor system, which performs the monitoring task, plays a vital role in the entire feedback loop. The placement of the sensors has a profound impact on the accuracy of the monitoring, thereby affecting the quality of the loop. Thermal monitoring is one concrete example of operating condition monitoring we are particularly interested in. We propose both non-uniform and uniform sensor placement

techniques and compare these two approaches on two different styles of layouts of chip multiprocessors. We first experiment with the thermal-aware K-Means algorithm (a non-uniform placement technique), which has been shown to be successful for single core sensor placement. Our results revealed that this method might not be suitable for some CMP architectures. We then analyzed a grid based sensor placement and observed that the number of sensors should be increased dramatically to achieve high accuracy. We improved upon this solution by introducing a) an interpolation scheme and b) a dynamic selection scheme. The interpolation scheme is able to reduce the average errors for a given distribution. On the other hand, the dynamic selection method provides reasonable accuracy, by using only a relatively small fraction of embedded sensors per core. Overall, these schemes achieve accurate temperature readings with small number of embedded sensor readings. For instance, our sensor placement techniques are able to cover 98.7% of all the thermal emergencies using only 16 sensors per core.

Self-Adjusting Pipeline (SAP) Architecture

Our proposed Self-Adjusting Pipeline (SAP) architecture employs delay sensors to predict the timing hazard in a pipeline and adjust its timing behavior via tuning the clock skew to prevent the timing violation in case timing hazards are predicted. Such dynamic timing adjustment makes it possible to intervene before the timing variation actually manifests itself as a violation. We formulate the delay sensor insertion and variable clock skew configuration problem as a stochastic mixed-integer programming problem and propose a simulated-annealing based algorithm to solve it. A comparison between the designs with and without the self-adjusting enhancement reveals that, we are able to improve the average performance of a batch of chips by 9.5%.

B. Local Feedback Loop Based Self-Adjusting Mechanism:

Local feedback loop based self-adjusting mechanism employs specially design circuit elements which can sense the environmental changes and adjust themselves to ensure that the entire system adapts to the changes. In the following, we will propose a Self-Adjusting Clock Tree Architecture (SACTA), which implements the idea of local feedback loop based self-adjusting architecture at circuit level.

Self-Adjusting Clock Tree Architecture (SACTA)

SACTA guarantees correct timing behavior in pipelined circuits within a large range of thermal conditions through a self-adjusting, temperature-sensitive skew distribution mechanism. SACTA exploits clock skews to “steal” time from adjacent pipeline stages to maintain the performance in the presence of delay variations within pipeline stages. The self-adjustable skews are generated by a set of special skew buffers. These elements are designed to exhibit carefully tuned, temperature dependent delay behavior in synchronization with the temperature levels prevalent in the logic of the pipeline. SACTA fit into the category of local feedback loop based self-adjusting mechanisms in that the specially designed skew buffers perform both the monitoring and adjusting tasks.

These special delay elements need to be carefully configured to ensure resilience of the entire circuit against temperature variation. To determine their configurations, we proposed an efficient and general clock tree design and optimization framework. Furthermore, we show that SACTA is applicable across a wide spectrum of circuits, including multi-V_{dd}/V_{th} designs. Experimental results show that a pipeline supported by SACTA is able to prevent thermal induced timing violations within a significantly larger range of operating temperatures (on average, the violation-free range can be enhanced by over 15°C).

II. LIST OF RELATED PUBLICATIONS

1. **J. Long,** S. Ogrenci Memik, G. Memik, and R. Mukherjee, “Thermal Monitoring Mechanisms for Chip Multiprocessors”, ACM Trans. on Architecture and Code Optimization (TACO), Vol. 5, No. 2, August 2008.
2. **J. Long**, and S. Ogrenci Memik, “Automated Design of Self-Adjusting Pipelines”, ACM/IEEE Design Automation Conference (DAC), June 8-13, 2008.
3. S. Ogrenci Memik, R. Mukherjee, M. Ni, and **J. Long,** “Optimizing Thermal Sensor Allocation for Microprocessors”, IEEE Trans. on Computer-Aided Design (TCAD), Vol. 27, No. 3, March 2008.
4. **J. Long**, J. Ku, S. Ogrenci Memik, and Y. Ismail, "A Self-Adjusting Clock Tree Architecture to Cope with Temperature Variations", ACM/IEEE Int. Conf. on Computer-Aided Design (ICCAD), November 5-8, 2007 [Best Paper Award Finalist].