

Temperature-Aware Processor Frequency Assignment for MPSoCs Using Convex Optimization

Srinivasan Murali^{†‡}, Almir Mutapcic[‡], David Atienza^{†+},
Rajesh Gupta[§], Stephen Boyd[‡], and Giovanni De Micheli[†]

[†]LSI, EPFL, Switzerland

[‡]Department of Electrical Engineering, Stanford University, USA

[§]Department of Computer Science and Engineering, UCSD, USA

⁺DACYA, Complutense University of Madrid (UCM), Spain

[†]{srinivasan.murali, david.atienza, giovanni.demicheli}@epfl.ch,

[‡]{almirm, boyd}@stanford.edu, [§]rgupta@ucsd.edu

ABSTRACT

The increasing processing capability of *Multi-Processor Systems-on-Chips (MPSoCs)* is leading to an increase in chip power dissipation, which in turn leads to significant increase in chip temperature. An important challenge facing the MPSoC designers is to achieve the highest performance system operation that satisfies the temperature and power consumption constraints. The frequency of operation of the different processors and the application workload assignment play a critical role in determining the performance, power consumption and temperature profile of the MPSoC. In this paper, we propose novel convex optimization based methods that solve this important problem of temperature-aware processor frequency assignment, such that the total system performance is maximized and the temperature and power constraints are met. We perform experiments on several realistic SoC benchmarks using a cycle-accurate FPGA-based thermal emulation platform, which show that the systems designed using our methods meet the temperature and power consumption requirements at all time instances, while achieving maximum performance.

Categories and Subject Descriptors

B.8 [Performance and Reliability]: Performance Analysis and Design Aids

General Terms

Design, Performance

Keywords

Temperature-aware, thermal, MPSoCs, convex optimization

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CODES+ISSS'07, September 30–October 3, 2007, Salzburg, Austria.
Copyright 2007 ACM 978-1-59593-824-4/07/0009 ...\$5.00.

1. INTRODUCTION

Business analysts forecast a 200 billion dollar market for media-rich, mobile *System-on-Chip (SoC)* terminals in the near future [1]. These forthcoming SoCs will be composed in future technology nodes of multiple homogeneous and/or heterogeneous processors, namely *Multi-Processor System-on-Chip (MPSoC)* platforms.

However, the semiconductor industry is still facing several technological challenges to build these systems. In the near future, power dissipation and consequent thermal implications will be often the dominant restriction to MPSoC performance and a significant element in overall cost. In fact, prohibitively expensive packaging, heat sinks and additional cooling solutions at the physical level will be necessary to handle the processor temperatures in next technology nodes [2]. Localized temperature increase can produce transient reduction in overall system performance, unreliable timing delay variations or even permanent damages in the devices [3]. To achieve a reliable and efficient system operation, the MPSoCs need to operate below a maximum temperature value and power budget. Determining the operating frequencies of the different processors of the MPSoCs, such that the performance of the system is maximized, while satisfying the temperature and power budget constraints, is a challenging task.

In this paper, we present novel design methods based on convex optimization to solve this critical problem of determining the operating frequencies of the processors, such that the system performance is maximized and the temperature and power budget constraints are satisfied at all time instances of operation. We integrate accurate thermal models of MPSoCs, which have been validated against 3D finite-element analysis for estimating the heat flow of the chip (presented in [29]), in a mathematical optimization framework. We present methods that can be applied to two different types of systems: one that has support for Dynamic Voltage and Frequency Scaling (DVFS), where each processor's frequency can be varied over time, and the other where no DVFS is applied. As both types of MPSoCs exist today, we address the design of both systems.

We validate the solutions produced by the methods by integrating them onto a cycle-accurate FPGA-based thermal emulation platform [28] and perform experiments on several realistic SoC benchmarks. We extract the real temperature measurements from the platform, which show that the system always operates below the maximum temperature constraint. We also present ways in which our methods can be used to perform design space exploration stud-

ies, such as detailed trade-offs between performance and temperature constraints for different working conditions, which can be used very early in the system design flow.

We have the following realistic assumptions for the design process: (1) We assume that the goal of our optimization methods is to maximize the total amount of work done by the processors. During system design, this problem is commonly encountered in practice. As an example, when designing chip multi-processors, an important performance goal is to maximize the total amount of instructions that can be executed in certain time [27]. This workload maximization objective translates to maximizing the sum of the frequencies of the different processors over time, as a faster processor can execute more instructions in a given time interval. We also solve a problem variation, where the application workload is fixed, and the objective is to minimize system power consumption. (2) Once the processor frequencies are set based on our methods, we assume that the actual application tasks are mapped onto the processors with compiler or OS support, as done widely in MPSoCs that have multiple processing elements [27].

2. RELATED WORK

The performance-power optimization of processors using DVFS policies has received considerable attention [4]-[6]. However, most of these works do not consider temperature constraints of the system. In the temperature-aware design domain, two important sub-problems have been considered: modeling the thermal behavior of the system and designing methods to control processor performance and power consumption to meet the temperature constraints. Several recent works have addressed the issue of on-chip thermal modeling [8]-[16]. At the physical level, various methods can be used to model the heat transfer in the substrate. Finite-difference time domain [8], finite element [10], model reduction [11], random walk [12, 13] and Green-function [14] based algorithms have been applied for on-chip thermal analysis. At the architectural level, [15] presents a thermal/power model for super-scalar architectures. The work presented in [16] investigates the impact of temperature and voltage variations across the die of embedded cores.

Based on these and other similar models, *Dynamic Thermal Management (DTM)* techniques have been suggested [17]-[20]. In [17], the use of the feedback control theory is proposed as a way to implement adaptive techniques in the processor architecture. In [18], a predictive frame-based DTM algorithm is presented. The authors in [20] perform extensive studies on empirical DPM techniques for thermal management. Their results show that DVFS can be very inefficient if the invocation time is not set appropriately. In [21]-[24], temperature-aware floorplanning is used to place circuit blocks, such that an even thermal profile is obtained. System-level solutions have been defined to reduce the temperature in MPSoCs using different scheduling mechanisms [25, 26]. Finally, in [27], a detailed review of thermal management techniques for multi-core architectures is presented.

Most of these existing thermal management techniques are based on monitoring and tuning of processor frequencies that do not result in optimal solutions. Moreover, they do not provide a guarantee that the temperature constraints will be satisfied at all instances of operation, which is critical for achieving system reliability.

3. TEMPERATURE MODELING

We consider two layers for heat flow modeling: silicon layer and the heat spreading copper layer. We divide the chip floorplan into several thermal cells of cubic shape. Each element of the design can be represented by one or more thermal cells of the silicon layer.

Thermal modeling is achieved by considering the heat conductances and capacitances of the cells [15], [29], and the model is similar to first-order RC circuits. In particular, we consider the model presented in Figure 1. Each cell in a layer is adjacent to a maximum of 4 other neighboring cells of the layer and a cell in the silicon layer is also adjacent to the corresponding top level cell of the copper layer. The copper layer cells have a thermal conductance to the ambient, which is at the ambient temperature t_{amb} . This thermal conductance models the heat flow by convection from the copper layer to the environment. For the computation of the different thermal conductances, we refer the reader to [15].

Let n be the number of processors, also equal to the number of thermal cells in each of the layers in the design. Thus, the total number of thermal cells in both the copper and silicon layers is $2n$. We denote the power consumption of processor i at time instance k with $p_{i,k}$, where $i = 1, \dots, n$. We also define $p_{i,k}$, where $i = n + 1, \dots, 2n$, to model the interaction of the copper layer cell i with the ambient at time instance k , as presented in [29].

The differential equation modeling the heat flow is given by:

$$C\dot{t}_k = -G(t_k)t_k + p_k, \quad k = 1, \dots, m, \quad (1)$$

where t_k is the temperature vector, with each entry $t_{i,k}$ modeling the temperature of thermal cell i at time instance k . The total number of time-steps used for temperature modeling is m . The matrix C is diagonal, with the entries representing the thermal capacitance of the cells (in Joules/Kelvin) and G is the thermal conductance matrix (the conductance values have units of Watts/Kelvin). The silicon thermal conductivity varies with temperature, in an approximately linear fashion [7] (shown in Figure 2). Hence matrix G is a function of temperature t_k .

The Equation (1) can be integrated numerically, giving:

$$t_{k+1} = A(t_k)t_k + Bp_k, \quad k = 1, \dots, m, \quad (2)$$

where

$$A(t_k) = (I - \delta C^{-1}G(t_k)) \quad \text{and} \quad B = \delta C^{-1}.$$

Here I is the identity matrix and δ is the actual time interval between any two time-steps, k and $k + 1$. The above thermal model has been validated against 3D finite element analysis and simulations in [29], where it has been shown that the model accurately predict the temperature variations.

We also solve for the steady state temperature vector t_{ss} , *i.e.*, the temperature distribution across the thermal cells when we achieve thermal equilibrium. At steady state, the differential of the temperature with time is zero, $\dot{t}_k = 0$. Then, Equation (1) can be written as:

$$t_{ss} = G(t_{ss})^{-1}p. \quad (3)$$

Here we drop the time subscript k from both vectors t and p since we are at steady state, and their values do not change with time.

4. PROBLEM FORMULATION

We consider two problems of interest: a steady-state case and a dynamic-state case.

Problem Statement: *The objective of our optimization problem is to set the frequency of operation of the processors, such that the total workload supported by the system (i.e. the total number of instructions executed in certain time) is maximized. For the steady-state case, we need to assign a single frequency to each processor in the design, such that in steady state, the temperature and*

¹In the rest of this paper, unless otherwise noted, we assume that the subscripts i and k are defined for the ranges $\{1, \dots, 2n\}$ and $\{1, \dots, m\}$, respectively.




Figure 1: Example cells and thermal conductances

power consumption values of the processors are below user-defined thresholds. For the dynamic-state case, the frequency of operation of the processors can be varied over time, such that the temperature and power consumption values at all time instances are below user-defined thresholds.

In the steady-state formulation, each processor is assigned a constant frequency and voltage at all time instances. However, many MPSoCs have the capability of utilizing DVFS, where the frequency and voltage of operation of the processors can be varied over time, so that the system performance, power consumption and temperature distribution can be highly optimized. The dynamic-state formulation can be applied to systems that support DVFS, and the frequency and voltage of the processors are varied over time to better optimize the system performance.

The frequency of operation of the processors is represented by the vector f_k , with each entry $f_{i,k}$, $i = 1, \dots, n$, representing the frequency of processor i at time instance k . We seek to determine the optimal f_k values that maximizes workload under various constraints related to temperature and power consumption. An alternate formulation of the problem, where we seek to determine the optimal f_k values that minimizes power consumption for a given application workload (fixed number of instructions to be executed in certain time) is also presented. Technology and hardware restrictions imply a minimum (f_{\min}) and maximum operating frequency (f_{\max}) for the processors.

The operating voltage of a processor depends on the operating frequency, and this dependence varies with different process and technology generations. In this work we assume that the square of the voltage scales linearly with the frequency of operation, as it is a common method to scale voltage [30]. It is well-known that the dynamic power consumption of a CMOS circuit depends linearly on the frequency of operation of the circuit, and quadratically on the operating voltage [6]. Hence, the power consumption values at different time-instances can be obtained by quadratically scaling the power consumption of the processors at f_{\max} , *i.e.*,

$$p_{i,k} = p_{\max} f_{i,k}^2 / f_{\max}^2, \quad i = 1, \dots, n, \quad k = 1, \dots, m. \quad (4)$$

Here p_{\max} is the power consumption of the processors at the maximum frequency of operation, which can be obtained from the data sheet of the processors or from power simulations.

The constraints in our problem are that temperature and power consumption values of the processors are below user-defined thresholds. We denote the maximum safe operating temperature of the chip with t_{\max} and power budget for the design with p_{tot} .

We can now present the steady-state and the dynamic-state optimization problems. The goal of the *steady-state optimization problem* is to maximize the total workload that is supported by the system at steady-state by maximizing the sum of the frequencies of the




Figure 2: Silicon thermal conductivity and linear fit

different processors, while satisfying the given temperature constraints and power budget, *i.e.*,

$$\begin{aligned} & \text{maximize} && \mathbf{1}^T f \\ & \text{subject to} && G(t_{ss})^{-1} p \leq t_{\max} \\ & && p_{\max} f_i^2 / f_{\max}^2 = p_i, \quad i = 1, \dots, n \\ & && \mathbf{1}^T p \leq p_{\text{tot}} \\ & && f_{\min} \leq f \leq f_{\max}, \end{aligned} \quad (5)$$

where the optimization variables are vectors f and p , and $\mathbf{1}$ is the vector of all ones. Note that here we drop the time subscript k from f and p since in steady-state each processor is assigned a constant frequency and power at all time instances.

The goal of the *dynamic-state optimization problem* is to maximize the total workload, while satisfying the given temperature constraints and power budget at all time instances, *i.e.*,

$$\begin{aligned} & \text{maximize} && \sum_{k=1}^m \mathbf{1}^T f_k \\ & \text{subject to} && t_0 = t_{\text{amb}} \\ & && t_{k+1} = A(t_k)t_k + Bp_k, \quad k = 1, \dots, m \\ & && t_k \leq t_{\max}, \quad k = 1, \dots, m \\ & && p_{\max} f_{i,k}^2 / f_{\max}^2 = p_{i,k}, \quad i = 1, \dots, n, \quad \forall k \\ & && \sum_{k=1}^m \mathbf{1}^T p_k \leq p_{\text{tot}} \\ & && f_{\min} \leq f_k \leq f_{\max}, \quad k = 1, \dots, m. \end{aligned} \quad (6)$$

The problems (5) and (6) are non-convex optimization problems, and in general, such problems cannot be solved globally in polynomial time. On the other hand, an optimization model with convex constraints can be solved globally with polynomial time complexity (polynomial in the number of variables and constraints) [31]. There are two issues that make these two problems non-convex. The first is that the thermal equations (2) and (3) are non-convex, since the matrices A and G depend on temperatures in a complicated manner. The second issue is that the power equation (4) defines a non-convex (quadratic equality) constraint for both problems.

5. CONVEX OPTIMIZATION SOLUTIONS

In this section, we present novel optimization methods to solve the processor frequency assignment for the steady-state and the dynamic-state cases. We also show how to handle the non-convex thermal and power constraints using a 2-phase iterative algorithm.

5.1 Steady-State Optimization

We present the following algorithm for solving the problem (5).

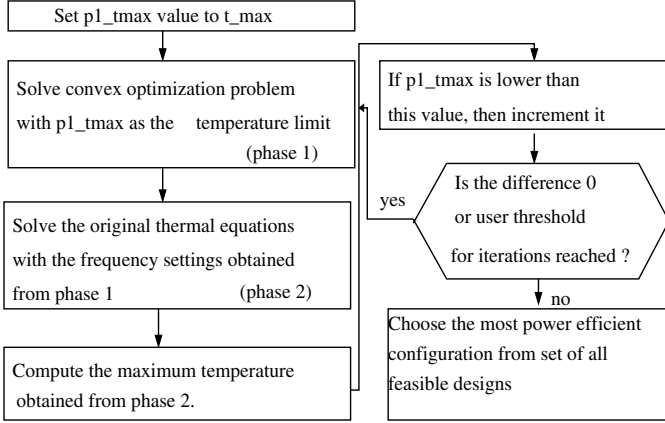
In the first phase (steps 1 and 2 of the algorithm), we take the matrix G to be constant in time, with thermal conductances computed at some given temperature. We observe from Figure 2 that the thermal conductivity decreases when the temperature increases. Thus, if we use the thermal conductivity value at t_{\max} , our thermal

Algorithm 1 Steady-state optimization

- 1: Set $G = G(t_{\max})$.
- 2: Solve the following convex optimization problem:

$$\begin{aligned}
 & \text{maximize} && \mathbf{1}^T f \\
 & \text{subject to} && G^{-1}p \leq t_{\max} \\
 & && p_{\max} f_i^2 / f_{\max}^2 \leq p_i, \quad i = 1, \dots, n \quad (7) \\
 & && \mathbf{1}^T p \leq p_{\text{tot}} \\
 & && f_{\min} \leq f \leq f_{\max},
 \end{aligned}$$

- 3: Let the resulting steady-state temperature vector be t_{opt} .
 - 4: Repeat steps 2 and 3, setting $G = G(t_{\text{opt}})$, until the temperature values across different iterations converge.
-


Figure 3: The 2-phase approach for the dynamic case

model will under-estimate the amount of heat flow from the processors when compared to the actual values. Consequently, the model would over-estimate the temperatures at different time instances when compared to the real temperature values. Therefore, a more conservative frequency assignment to the processors would be obtained, such that the maximum temperature allowed by the model is in reality lower than t_{\max} .

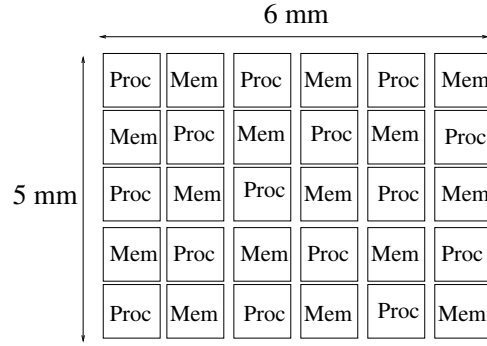
In problem (7), we also relax the power equation (4) to an inequality. It can be shown that the resulting relaxed convex problem is equivalent to the original problem with the quadratic equality constraint [31, pg. 191].

In the second phase (in step 4), we try to close the gap between the approximated thermal models and the real temperature values to achieve the optimum frequency values. Once the set of steady-state temperatures are obtained, we calculate new values of matrix G at the obtained temperatures. Then, with the $G(t_{\text{opt}})$ values, we re-run the first phase of the optimization. We repeat the two phases until the temperature results converge. An analysis of the convergence of the results from the two phases for an application is presented later (in Section 6.3).

5.2 Dynamic-State Problem Optimization

The algorithm used to solve the dynamic state workload assignment problem is presented in Figure 3. For the dynamic state problem, the temperature values obtained from phase 1 can be different for different time instances. Thus, the second phase of the steady-state formulation cannot be directly applied. To this end, once the processor frequencies are obtained, in the second phase we re-run the original thermal equations (Equation 2) and obtain the observed maximum temperature values. Note that once the frequency points are obtained, the original thermal equations can be readily solved.

For the dynamic state problem, we solve the following relaxed


Figure 4: Floorplan of multi-media SoC with 30 cores

convex optimization problem:

$$\begin{aligned}
 & \text{maximize} && \sum_{k=1}^m \mathbf{1}^T f_k \\
 & \text{subject to} && t_0 = t_{\text{amb}} \\
 & && t_{k+1} = A(p1_{t_{\max}})t_k + Bp_k, \quad k = 1, \dots, m \\
 & && t_k \leq p1_{t_{\max}}, \quad k = 1, \dots, m \\
 & && p_{\max} f_{i,k}^2 / f_{\max}^2 \leq p_{i,k}, \quad i = 1, \dots, n, \forall k \\
 & && \sum_{k=1}^m \mathbf{1}^T p_k \leq p_{\text{tot}} \\
 & && f_{\min} \leq f_k \leq f_{\max}, \quad k = 1, \dots, m,
 \end{aligned} \quad (8)$$

where $p1_{t_{\max}}$ is initially set to t_{\max} .

After running the second phase, we obtain the differences between the maximum temperature values obtained from the approximated thermal model used in phase 1 and the original thermal model (as shown in Figure 3). Then, we can increase the t_{\max} value that is fed to the optimization model in phase 1 (represented by the $p1_{t_{\max}}$ value in Figure 3). After this, we re-run both phases, increasing the $p1_{t_{\max}}$ value, until the maximum value obtained from phase 1 converges to the value from phase 2 or until the number of repetitions reaches a user-defined threshold.

5.3 Problem Variations

Our optimization methodology can be used to model several variations of the temperature-aware workload assignment problem. As an example, in this sub-section, we consider an alternate problem of setting the frequencies of the processors such that a pre-defined amount of workload is executed with minimum power consumption, satisfying the temperature constraints.

Let w be the total workload, in terms of the number of processor cycles, to be assigned to the processors. Let l be the total time by which the workload needs to be completed by the different processors. Then, the average frequency of operation of each of the processor that would be required to complete the workload is given by $f_{\text{avg}} = w / (n \times l)$.

As an example, an application may require six million processor cycles of computation, which needs to be completed in 10 ms. With 6 processors, each processor on average needs to operate at 100 MHz to execute the application in the required time interval.

The optimization problem in Equation (8) is solved using the following objective function:

$$\text{minimize} \quad \sum_{k=1}^m \mathbf{1}^T p_k \quad (9)$$

and with the following additional constraint used along with the other constraints in Equation (8):

$$\sum_{k=1}^m \mathbf{1}^T f_k \geq m \times n \times f_{\text{avg}}. \quad (10)$$




Figure 5: Normalized Workload




Figure 6: Frequency variation for multi-media benchmark




Figure 7: Temperature variation for multi-media benchmark

6. EXPERIMENTAL RESULTS

We have applied the proposed convex-based thermal-aware DVFS optimization approach to several case studies to show its effectiveness in complex MPSoC designs.

6.1 Temperature Constraint Effects

In the first set of experiments, we have applied our optimization technique on a relatively large MPSoC platform implementing a multimedia benchmark suite. The studied MPSoC consists of 30 cores [32]: 15 processors and 15 on-chip memories that are evenly distributed around the chip as shown in Figure 4. The processors are ARM-based cores, which closely fit in the dimension of $1\text{mm} \times 1\text{mm}$ used in our floorplan, typical of today’s embedded processors [33]. The on-chip memories have sizes that also fit well in the same dimensions. The platform simulates multi-media applications, such as video texture coding and 3-D reconstruction. The power consumption values are obtained by accurately estimating the instruction level power dissipation in the functional units, based on switching activities of the components, accurately extracted on a Virtex-II Pro based emulation platform [28].

In order to see the effect of temperature on the design, we vary the maximum temperature that can be tolerated by the chip and find the amount of workload that can be supported by the design. These values are obtained using our proposed optimization methods, which are solved using CVX [34], a convex optimization solver. Such an analysis is very useful for designers to perform early design trade-offs studies. As an example, when an expensive packaging solution is utilized or a more temperature-tolerant circuit design is used, the maximum temperature that can be tolerated by the chip can be higher. However, it is a challenging task for the designer to estimate the potential gains of such approaches early in the design cycle, but such a study can be easily and accurately obtained by the presented methods. The resulting design points for the considered multi-media design are shown in Figure 5. In this set of experiments, we set $t_{\text{amb}} = 300\text{K}$. From the figure, we see that at low values of maximum temperature threshold, the normalized workload curve (workloads are normalized with the workload at 403K) is slightly steeper than at higher temperatures. This is because, the thermal conductivity of silicon decreases with increase in temperature (see Figure 2); hence the relative performance improvement reduces with temperature increase.

Next, the temperature and frequency variations of one of the processors over time (with the maximum temperature set to 373K) are presented in Figures 6 and 7.

When the system starts its operation, the chip is at the ambient temperature. Hence, the frequency of operation and the power dissipation of the processors can be high without any thermal risk.

However, as soon as the chip starts heating up, the maximum temperature constraint forces a drop in the frequency of operation.

6.2 Effect of hot-spots

In many MPSoCs, along with several processing elements, other cores are also present, such as dedicated hardware blocks, memories, *etc.* In such a setting, when the workload is split across the different processor cores, the thermal effects of the adjoining cores play a vital role in the workload assignment pattern. A good example is the new IBM cell architecture [35]. In the design (the floorplan is shown in Figure 8), there are multiple *Synergistic Processing Elements (SPEs)* that are RISC processors, to which the workloads are assigned by a *Power Processor Element (PPE)*. In [35], the authors show that the PPE has a thermal hot spot. Here we study the impact of workload assignment to the different SPEs using our method, considering the thermal hot-spot of the PPE.

The frequency assignment achieved by our method and the resulting temperature profiles of 4 different SPEs are presented in Figures 9 and 10. In the figures, we present the results for the two SPEs near the hot-spot of the PPE (SPE 1 and SPE 5) and for the two SPEs that are far away from the hot-spot (SPE 4 and SPE 8). The results clearly show that the workload assignment is highly uneven, with the thermally-efficient SPEs supporting a higher load. Intuitively, though a designer could easily observe that a thermally efficient SPE would support a higher load, it is very difficult to estimate how to split the workload across the SPEs in an efficient manner. Our methodology exactly addresses this problem and can be utilized as a powerful and useful tool to obtain efficient results and perform detailed studies of different assignment policies.

6.3 Convergence Analysis

In this sub-section, we show the convergence between the different iterations of the steady-state optimization algorithm (presented in Section 5.1) for the multi-media MPSoC platform. In Figure 11, the minimum temperature values (across all the thermal cells) that are obtained after running each iteration are presented. As we approximate the thermal conductances, in the first iteration, the models result in a conservative frequency assignment, leading to lower actual temperature values. As we close the gap between the predicted and actual values in the next iterations, the frequency assignment and the temperature values gradually stabilize.

For all our experiments, the run-time of the convex optimization models were less than few hours, when run on a 3.2 GHz PC with 1 GB RAM. This is due to the fact that the optimization problems are solvable in polynomial time and are highly scalable to large problem instances.




Figure 8: Cell processor floorplan




Figure 9: Assigned frequencies




Figure 10: Temperature values




Figure 11: Convergence results of steady-state formulation

7. CONCLUSIONS

Forthcoming MPSoCs can provide a very high processing power for embedded devices. However, one of the major design challenges is that the power dissipation of MPSoCs is rapidly increasing, which leads to large increase in chip temperature. This can severely affect system performance, and even produce permanent damages in the device in extreme cases. Achieving a high-performance operation, satisfying temperature and power constraints is critical for MPSoCs. In this paper, we have presented novel design approaches to perform optimal thermal-aware frequency assignment in MPSoCs, based on convex optimization. Our methods guarantee that the chip temperature remains below the maximum threshold at all time instances of operation. Our experimental results show that, even in very complex MPSoCs, our approach enables to setup optimal frequencies for each processor, as well as provides designers with a consistent framework to perform studies of trade-offs between performance and temperature very early in the design flow.

8. ACKNOWLEDGMENTS

This research is supported by a grant from the US National Science Foundation (under contract CCR-0305718) and a grant from the Fonds National Suisse (FNS, Grant 20021-109450/1).

9. REFERENCES

- [1] M. Viredaz and D. Wallacha, "Power evaluation of a handheld computer", *IEEE Micro*, January 2003.
- [2] S. Borkar, "Design challenges of technology scaling", *IEEE Micro*, July-Aug 1999.
- [3] O. Semenov et al., "Impact of self-heating effect on long-term reliability and performance degradation in CMOS circuits", *IEEE Transactions on Devices and Materials*, March 2006.
- [4] H. Aydin et al., "Dynamic and aggressive scheduling techniques for power-aware real-time systems", *Proc. IEEE Real-Time Systems Symposium*, pp. 95-105, Dec. 2001.

- [5] Y. Shin et al., "Power optimization of real-time embedded systems on variable speed processors", *Proc. ICCAD*, 2000.
- [6] R. Jejurikar, R. Gupta, "Optimized Slowdown in Real-Time Task Systems", *Proc. ECRTS*, 2004.
- [7] A. Ramalingam et al., "Accurate Thermal Analysis Considering Nonlinear Thermal Conductivity", *Proc. ISQED* 2006.
- [8] T.-Y. Wang and C.-P. Chen, "3-d thermal-adi: A linear-time chip level transient thermal simulator," *IEEE TCAD*, December 2002.
- [9] J. Deeney, "Thermal modeling and measurement of large high power silicon devices with asymmetric power distribution," *Proc. of the International Symposium on Microelectronics*, 2002.
- [10] B. Goplen and S. Sapatnekar, "Efficient thermal placement of standard cells in 3d ics using a force directed approach", *Proc. ICCAD* 2003.
- [11] T. Wang and C. Chen, "Spice-compatible thermal simulation with lumped circuit modeling for thermal reliability analysis based on model order reduction," 2004.
- [12] H. Qian et al., "Random walks in a supply network," in *Proceedings of the 40th Design Automation Conference (DAC)*. ACM Press, 2003, pp. 93-98.
- [13] H. Qian and S. Sapatnekar, "Hierarchical random-walk algorithms for power grid analysis in a supply network", *Proc. ASPDAC*, 2003.
- [14] Y. Zhan and S. Sapatnekar, "Fast computation of the temperature distribution in vlsi chips using the discrete cosine transform and table look-up", *ASPDAC* 2005.
- [15] K. Skadron et al., "Temperature-aware microarchitecture: Modeling and implementation", *TACO*, 2004.
- [16] H. Su et al., "Full chip leakage estimation considering power supply and temperature variations", *Proc. ISLPED*, Aug. 2003.
- [17] K. Skadron et al., "Thermal-rc modeling for accurate and localized dynamic thermal management", *Proc. IEEE HPCA*, 2002.
- [18] J. Srinivasan and S. V. Adve, "Predictive dynamic thermal management for multimedia applications", *ICSO3*, June 2003.
- [19] M. Huang et al., "A framework for dynamic energy efficiency and temperature management", *MICRO*, 2000.
- [20] D. Brooks and M. Martonosi, "Dynamic thermal management for high-performance microprocessors," *HPCA*, 2001.
- [21] C. Chu and D. Wong, "A matrix synthesis approach to thermal placement," *Proc. TCAD*, Nov 1998.
- [22] C. Tsai and S. Kang, "Cell-level placement for improving substrate thermal distribution", *Proc. TCAD*, Feb 2000.
- [23] G. Chen and S. Sapatnekar, "Partition-driven standard cell thermal placement," *ISPD* 2003.
- [24] B. Goplen and S. Sapatnekar, "Thermal via placement in 3d ics," *ISPD*, 2005.
- [25] E. Rohou and M. Smith, "Dynamically managing processor temperature and power", *Proc. of 2th workshop on Feedback-directed optimization*, November 1999.
- [26] W. Hung et al., "Thermal-aware allocation and scheduling for systems-on-chip", *DATE*, 2005.
- [27] J. Donald, M. Martonosi, "Techniques for Multicore Thermal Management: Classification and New Exploration", *Proc. ISCA*, pp. 77-88, 2006.
- [28] D. Atienza et al., "A Fast HW/SW FPGA-Based Thermal Emulation Framework for Multi-Processor System-on-Chip", *Proc. DAC*, pp. 618-623, 2006.
- [29] G. Paci et al., "Exploring temperature-aware design in low-power MPSoCs", *Proc. DATE*, pp. 838-843, 2006.
- [30] S. Murali, et al., "Mapping and configuration methods for multi-use-case networks on chips", *Proc. of ASP-DAC*, 2006.
- [31] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [32] F. Angiolini, et al., "Contrasting a NoC and a traditional interconnect fabric with layout awareness", *Proc. of DATE*, 2006.
- [33] S. B. Furber, *ARM System-on-Chip Architecture*, Addison-Wesley Longman Publishing Co., 2000.
- [34] M. Grant, S. Boyd, and Y. Ye. *CVX: Matlab software for disciplined convex programming*, version 1.0 beta 3. Available at www.stanford.edu/~boyd/cvx/.
- [35] O. Takahashi, et al., "Power-conscious design of the cell processor's synergistic processor element." *IEEE Micro*, 2005.