

Temperature Sensor Placement in Thermal Management Systems for MPSoCs

Francesco Zanini[†], David Atienza^{*}, Colin N. Jones[‡], Giovanni De Micheli[†]

[†] Laboratory of Integrated Systems (LSI), EPFL, Switzerland

^{*} Embedded Systems Laboratory (ESL), EPFL, Switzerland

[‡] Automatic Control Laboratory, ETH Zurich, Switzerland

e-mail: {name.surname}@epfl.ch, cjones@control.ee.ethz.ch

Abstract—Modern high-performance processors employ thermal management systems, which rely on accurate readings of on-die thermal sensors. Systematic tools for analysis and determination of best allocation and placement of thermal sensors is therefore a highly relevant problem. This paper proposes a novel technique for determining the placement of temperature sensors on complex *Multi-Processor Systems-on-Chips* (MPSoCs) floorplans. The proposed method first analyzes the observability of the system for all the possible sensor placement configurations. Minimum sensors placements ensuring the observability of the portion of the MPSoC system that is relevant to the designer are then compared with simulation-based data coming from a wide set of benchmarks. Pareto points identifying the best configurations are then stored. According to user designer needs the best configuration is selected and a specific location is assigned to each sensor. We compared the proposed method with state-of-the-art approaches [5], [6]. Results show a reduction up to $4.5\times$ in the number of required sensors.

I. INTRODUCTION

The number of functional units and cores integrated on a chip is increasing. Examples include the Sun's Niagara [2] and the Tiler's 64-core architecture [3]. This increase in capability of MPSoCs is leading to an increase in chip power dissipation, which in turn leads to significant increase in chip temperature. Temperature gradients and hot-spots not only affect the performance of the system, but also lead to unreliable circuit operation and affect the life-time of the chip [4]. Meeting temperature constraints and reducing hot-spots is indeed an important and difficult challenge facing the MPSoC designers. Thus thermal management approaches have been proposed [7]- [12], but all these techniques require on-line thermal profile information from the chip to perform frequency assignment optimization.

In this work we focus on the thermal sensors placement problem. Our goal is to minimize the number of sensors while maximizing the thermal profile estimation accuracy.

We propose a novel approach for choosing a sensor configuration based on observability analysis. This is a general approach that can be potentially applied to any MPSoC architecture. In this work, we validated the system using a commercial MPSoC from SUN. We compare the resulting sensor placement with state-of-the-art algorithms. Results show a reduction up to $4.5\times$ in the number of required sensors.

II. RELATED WORK

Temperature management at system-level has been presented in [7]- [12]. Several groups have addressed the problem of thermal modelling and simulation at different levels of abstraction. Finite-difference time domain [14], [15] based algorithms have been proposed. In [16] a thermal/power model for super-scalar architectures is presented. One problem related to all these techniques is that they require on-line thermal profile information from the chip in order to perform the frequency and voltage assignment optimization.

A study of the thermal profile estimation problem has been analyzed in [5] and [6]. The proposed solutions are based on techniques trying to reduce temperature differences between thermal sensors and hot-spots by using the minimum number possible of sensors for a certain accuracy. The problem with these approaches is that since hot-spots are application dependent, there is no guarantee that all hot-spots are detected during the lifetime of the device.

In [17] the authors select the location of the sensing element according to a gramian-based sensor strategy. In [18] the problem of making a system observable is solved by employing of graph theory. The problem of choosing a set of measurements from a much larger set that also minimizes the estimation error is solved by [19] using a convex optimization based approach. This last method approximately solves the problem and has no guarantee that the performance gap is always small.

III. PROPOSED METHOD

A. Thermal model

To model the thermal properties of the MPSoC, we use the finite element model presented in [16] and [15]. The model is composed by two types of layers: the silicon layer and the heat-spreading copper layer. The chip floorplan is divided into several thermal cells of cubic shape. Every single functional unit in the floorplan can be represented by one or more thermal cells of the silicon layer. Thermal modelling is computed considering the heat conductances and capacitances of the cells as computed and validated in [16] and [15]. The differential equations modelling the heat flow are given by solving this network. Many methods and ODE solvers can be used. A survey on these methods is reported in [13].

The thermal model that we want to represent is slightly nonlinear since coefficients are temperature-dependent (relative error in the order of 0.16%) [15]. To represent the thermal

model using a linear, time invariant discrete-time system, the solution of the differential equations modelling the heat flow inside the MPSoC has been linearized. The rationale behind it is described in [13], [16] and [15]. In the sequel we assume that the k^{th} temperature measurement is done at time t_k . The system can be represented with the equations:

$$x(k+1) = Ax(k) + Bu(k) + w \quad (1)$$

$$y(k) = Cx(k) \quad (2)$$

where at time k , $x(k)$ is the plant's state, $u(k)$ is its input and $y(k)$ is its output. The temperature value of each cell is the state $x \in \mathbb{R}^{2n}$ of our system. The first n entries represent the cells composing the silicon floorplan and the remaining n entries model the copper layer. The input of the system $u \in \mathbb{R}^p$ is related to the frequencies of the cores according to [16]. The output $y \in \mathbb{R}^s$ of our system is the temperature observed by the s on-chip thermal sensors placed in the silicon layer. Matrices A, B, C and vector w describe the system and model all geometric constraints among each entry of the state vector and its placement on the chip floorplan. Matrix $A \in \mathbb{R}^{2n \times 2n}$ expresses the part of the temperature spreading process inside the chip that depends only on the current temperature profile of the cells. Matrix $B \in \mathbb{R}^{2n \times p}$ expresses the temperature increase due to the input. The part of the system dynamic that is not controllable by the input vector such as the heat dissipation of the copper layer due to room temperature is expressed by vector $w \in \mathbb{R}^{2n}$. Matrix $C \in \mathbb{B}^{s \times 2n}$, $\mathbb{B} = \{0, 1\}$, represents a selection matrix that models the placement of a sensor on the silicon die. Namely $c_{i,j}$ is equal to 1 if thermal sensor i is located inside the cell j . Since we are assuming to have distinct measurements coming from distinct sensors, C has only 1 nonzero element per row. For technological reasons thermal sensors can be placed only on the silicon layer, so sensors can be located only for $i, j \leq n$.

In the model described by Equations 1 and 2, the sampling frequency $f = (t_{k+1} - t_k)^{-1}$ is the frequency at which we are assuming to sample the original continuous time system. The higher this frequency is, the higher the information thermal sensors are providing about the MPSoC thermal spreading process. The contribution of this paper is a methodology to find sensor placements leading to Pareto points in the area versus frequency plane.

B. Observability and Sensor Placement

Observability refers to the property of a system that enables the reconstruction of the state variables given the inputs [1]. For the system identified by Equations 1 and 2, it means that we are able to reconstruct completely the thermal profile of the chip given the inputs only by looking at the measurements coming from the sensors, placed in locations specified by the matrix C . This means that we are assuming to have in the output vector s distinct temperature measurements coming from s distinct cells. The rank of the observability matrix Q expresses the number of states that can be reconstructed

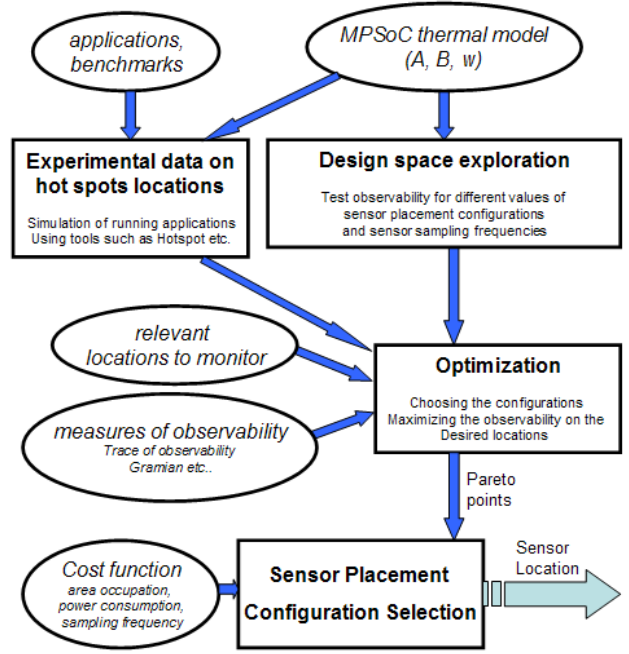


Fig. 1. Proposed method block diagram

from the measurement vector y . The observability matrix Q is expressed by the following equation (see [1]):

$$Q = [C; CA; \dots; CA^{2n-1}] \quad (3)$$

If the rank of Q equals $2n$, the state vector x can be reconstructed completely from the measurement vector y and the input vector u .

The problem of selecting the right placement of thermal sensors to both minimize the number of sensors and maximize observability is the problem of choosing the matrix C with the minimum number of rows that maximize the rank of the observability matrix Q . Given a determined MPSoC model, this problem depends on the location and the number of sensors inside floorplan (matrix C). The sensor sampling frequency (f) has as well an influence on the observability because matrices of Equations 1, 2 depend on f .

C. Algorithm

The block diagram of the proposed algorithm is presented in Figure 1. The method is consists of four steps.

In the first step, experimental data of hot spots locations are recorded during the runtime execution of the system. Data can be obtained from real chip temperature measurements or from simulations using tools such as Hotspot. These data will point out in which locations an accurate monitoring is needed in order to identify the rising of potential Hotspots. It's important to notice that these data are not used to define the placement for the sensors. They are used in step 3 of the algorithm as a criterion to rank among different sensor placements having the same observability properties.

In the second step the design space exploration is done on all the possible sensor placement configurations. First the model of the system of Equations 1,2 is sampled using a frequency

f that ranges from F_{min} to F_{max} . After that, the number of sensors employed in the placement is varied from 1 to n . A value of 1 means having only 1 sensor in the whole floorplan, while a value of n means having a sensor per cell in the silicon layer. At this stage for every value of s and f , a total of $\binom{2n}{s}$ sensor placement configurations are generated. The possible configurations have only one sensor per cell. This leads to have a matrix C having one nonzero element per row and a total of s rows. The rank of the observability matrix Q is computed for each configuration.

The third step performs an optimization based on data collected on both previous steps plus some additional data. This step performs a selection of the number of all analyzed sensor placement configurations. First, configurations that do not allow the estimation of area of the chip that are relevant to the designer are discarded. If a full profile estimation is needed by the designer, then, placements leading to observability matrixes with rank less than $2n$ are discarded. As a second criterion, configurations where sensors are placed on experimental hot-spots locations (see step 1 of the algorithm) are preferred to other ones. Remaining placements are ranked according to metrics to measure the observability of a system (i.e.the observability Gramian [17]- [19]). Finally according to aforementioned metrics, Pareto point placements are computed. A specific sensor placement is corresponding to every Pareto point in the plane sensor number s versus sensor sampling frequency f .

The last step selects the best placement according to the designer defined criteria based on area occupation(related to s), power consumption(related to s and f) and sensor sampling frequency (related to f).

IV. RESULTS

A. Experimental Setup

In our setup, we consider an architecture resembling the 8-core Niagara-1 (UltraSparc T1) architecture from Sun Microsystems [2]. The power consumption of all other elements has been chosen according to [2]. The floorplan of the Niagara-1 multicore architecture, is presented in Figure 2. Values regarding thermal resistance, silicon thickness and copper layer thickness have been derived from [2]. To simulate the system we use the execution characteristics of tasks from a mix of different benchmarks, ranging from web-accessing to playing multimedia [20]. The simulation step for the discrete time integration of the RC thermal model is $200\mu s$.

B. Placement Results

We applied the proposed algorithm to the case study described in the experimental setup. The overall computation of the proposed algorithm on a INTEL CoreTM2 duo laptop having a frequency of 2GHz (T7200) in the case of a modelling performed using 24 states took 3.44 minutes. After the first step, we obtain the design space exploration results of all possible sensor placement configurations. By plotting the percentage of the observable states over the overall states of the MPSoC versus the number of sensors used and their

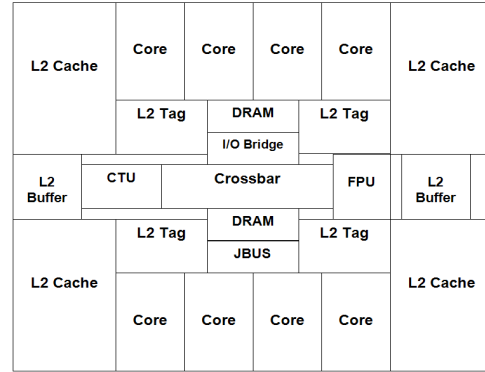


Fig. 2. Floorplan used of the Niagara-1 multicore case study

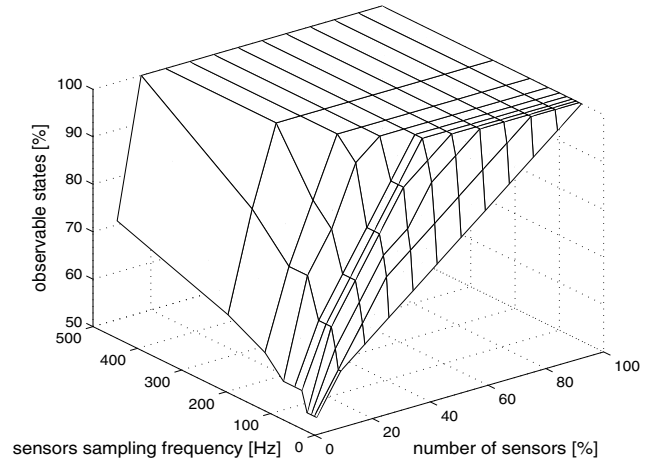


Fig. 3. Design space exploration of the case study (step 2)

sampling frequency, we obtain the plot of Figure 3. As it can be noted from the graph, for a fixed percentage of the observable states, there are many options. The graph shows that the number of sensors can be reduced by increasing the sensors sampling frequency and vice-versa. It is important to notice also that, there are many possible sensor placement configurations associated with any point in previous graph.

At this stage, among all possible placement configurations we identify Pareto points inside the design space. Trade-offs are between the number of sensors and their sampling frequency. The reason is because for a given observability target, the lower bound of thermal sensors employed depends on the thermal sensor sampling frequency (see Figure 3). Pareto points are computed by connecting previous results with experimental data derived from simulations. Additional data are the following. According to simulations and results from [20] and [2], in our case study, critical areas for hotspots are the cores and the crossbar located in the central part of the chip. Moreover, we are interested in monitoring 100% of the states of our MPSoC. In this case study, the overall Pareto points computation step takes few seconds. The graph identifying the resulting Pareto points in the plane sensors number Vs sensors sampling frequency is shown in Figure 4.

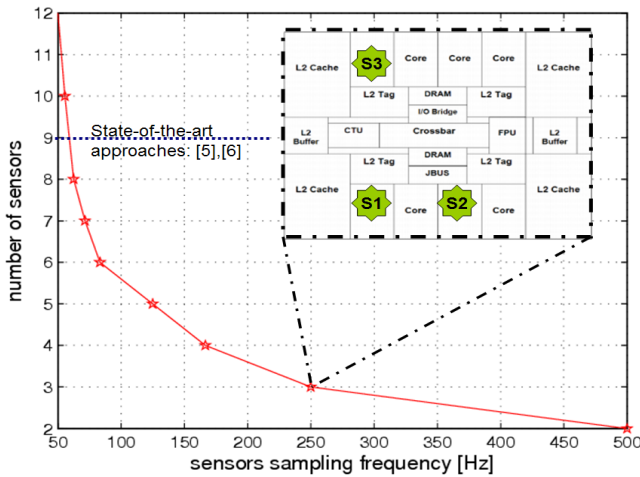


Fig. 4. Pareto points (steps 3+4) and comparison with [5],[6].

In the last step of the algorithm, we assume as possible design constraint to allow a maximum of 3 thermal sensors on the MPSoC. Moreover we want to have a sensor sampling frequency as low as possible. According to Figure 4, the corresponding Pareto point is a 3 sensor configuration with a sampling frequency of 250Hz . This means that if we want to make the system observable with only 3 sensors, we need to sample them every 4ms and we need to place them in a specific configuration. This specific placement is shown in Figure 4. This sensor configuration supports a complete estimation of the system and so a complete reconstruction of the thermal profile of the MPSoC. This operation can be implemented by using conventional estimation techniques (i.e. Kalman filter [1]). This will also correct for potential noise sources present in thermal sensors.

We compare the method with approaches [5] and [6]. Algorithm [5] finds the sensor placements that provides the best estimation accuracy for a certain number of sensors, according to experimental hot-spots locations. The algorithm is based on an interpolation technique based on experimental data derived from simulations. Algorithm [6] minimizes the temperature differences among the hot-spot temperature and the one detected by the thermal sensor. Both these algorithms are based on techniques trying to catch hot spots by using the minimum number possible of sensors for a certain accuracy.

Our method is not targeting hot spots but the observability of the system. Once the system is observable, hot spots are automatically detected. The reason is because the portion of the thermal profile of the chip that is relevant to the designer can be completely reconstructed from sensors measurements. The advantage is a strong reduction in the number of required sensors. In our case study, the cores may run with independent frequencies. This imply that hot spots are uncorrelated from core to core. Consequently, according to [5] and [6] at least one sensor per independent unit is required to monitor those elements and detect potential hot spots. Thus, a minimum of 9 sensors (8 sensors for the cores plus 1 sensor for the crossbar) are needed by both techniques to detect all possible hot spots. Conversely our method needs only 3 sensors, i.e. a reduction of

$3\times$. Moreover Figure 4 shows that the gain can reach $4.5\times$ at a sampling frequency of 500Hz . The proposed method enables the full thermal profile estimation of the MPSoC by using a number of sensors that is smaller than the number of the location of potential hot-spots formation regions.

V. CONCLUSION

We propose a methodology for determining the placement of temperature sensors on complex MPSoC floorplans. The proposed algorithm identifies the best configurations of sensors leading to Pareto points in the design space in terms of sensors sampling frequency versus number of required sensors. We compared the sensor placement with state-of-the-art algorithms such as [5] or [6] in the case of a commercial MPSoC device. Results show that the proposed method offers a reduction up to $4.5\times$ in the number of required sensors. Moreover the method constructs a maximally observable system for a given number of sensors.

ACKNOWLEDGMENT

This research has been partially funded by the Nano-Tera.ch NTF Project CMOSAIIC (ref. 123618), which is financed by the Swiss Confederation and scientifically evaluated by SNSF.

REFERENCES

- [1] G.F.Franklin, et al., *Digital Control of Dynamic Systems*, McGraw Hill.
- [2] P. Kongetira, K. Aingaran, and K. Olukotun, *Niagara: A 32-way multi-threaded SPARC processor*, IEEE Micro, March-April 2005.
- [3] Tiler Corporation, *Tiler's 64-core architecture*, www.tiler.com/products/processors.php, 2007.
- [4] 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.
- [5] S.O.Memik et al. *Optimizing Thermal Sensor Allocation for Microprocessors*. IEEE TCAD, 2008.
- [6] S.Sharifi et al. *An analytical model for the upper bound on temperature differences on a chip*, Proc. of GLSVLSI 2008.
- [7] R. Mukherjee, et al. *Physical aware frequency selection for dynamic thermal management in multi-core systems*, Proc. of ICCAD, 2006.
- [8] D. Brooks et al. *Dynamic thermal management for high-performance microprocessors*, HPCA, 2001.
- [9] C J. Hughes, et al. *Saving energy with architectural and frequency adaptations for multimedia applications*, Proc MICRO, 2001.
- [10] S.Murali et al. *Temperature Control of High Performance Multicore Platforms Using Convex Optimization*, Proc. of DATE 2008, Germany.
- [11] F.Zanini et al. *Multicore Thermal Management with Model Predictive Control*, ECCTD 2009.
- [12] F.Zanini et al. *A Control Theory Approach for Thermal Balancing of MPSoCs*, ASPDAC 2009.
- [13] F.Zanini et al. *Optimal Multi-Processor SoC Thermal Simulation via Adaptive Differential Equation Solvers*, VLSISoC 2009.
- [14] T.-Y. Wang et al., *3-d thermal-adi: A linear-time chip level transient thermal simulator*. IEEE TCAD, December 2002.
- [15] G. Paci et al., *Exploring temperature-aware design in low-power MP-SoCs*, Proc. of DATE, 2006.
- [16] K. Skadron et al., *Temperature-aware microarchitecture: Modeling and implementation*, TACO, 2004.
- [17] C.Sumana et al., *Optimal selection of sensors for state estimation in a reactive distillation process*, Journal of Process Control, 2009.
- [18] S.Joshi et al., *Sensor selection via convex optimization*, transaction on signal processing, 2009.
- [19] T.Boukhobza et al., *State and input observability recovering by additional sensor implementation: a graph theoretic approach*, Automatica 2009.
- [20] A.K.Coskun et al., *Temperature Aware Task Scheduling in MPSoCs*, Proc. of DATE, 2007.