

Optimization of the Bias Current Network for Accurate On-Chip Thermal Monitoring

Jieyi Long and Seda Ogrenci Memik

Dept. of EECS, Northwestern Univ., Evanston, IL 60208
{jlo198, seda}@eecs.northwestern.edu

Abstract—Microprocessor chips employ increasingly larger number of thermal sensing devices. These devices are networked by an underlying infrastructure, which provides bias currents to sensing devices and collects measurements. In this work, we address the optimization of the bias current distribution network utilized by the sensing devices. We show that the choice between two fundamental topologies (the 2-wire and the 4-wire measurement) for this network has a non-negligible impact on the precision of the monitoring system. We also show that the 4-wire measurement principle supports the remote sensing technique better. However, it requires more routing resources. We thus propose a novel routing algorithm to minimize its routing cost. We also present a detailed evaluation of the quality of the resulting system in presence of process and thermal variations. Our Monte Carlo simulations using the IBM 10SF 65nm SPICE models show that the monitoring accuracies can be as high as 0.6°C under considerable amount of process and temperature variation. Moreover, by adopting a customized routing approach for the current mirror network, the total wire length of the bias current network can be reduced by as much as 42.74% and by 27.65% on average.

I. INTRODUCTION

Precise runtime thermal monitoring plays a crucial role in maintaining the performance and reliability of high-performance microprocessors. Thermal sensors are widely used to guide thermal monitoring. There is a clear trend that more and more on-chip thermal sensors are incorporated into these systems. For instance, an early version of Intel® Pentium® M processor was equipped with two on-core thermal sensors [1], and later, in a 90nm Intel® Itanium® processor, four thermal sensors are placed on the chip [2]. More recently, AMD proposed a quad-core Opteron™ processor, where a total number of thirty-eight thermal sensors are deployed for precise thermal monitoring [3].

In order to maximize the coverage, the thermal sensing devices are scattered across the entire chip. They are networked by an underlying infrastructure, which provides the bias currents to the sensing devices, collects measurements, and performs analog to digital signal conversion. Therefore, the supporting infrastructure is an on-chip element at a global scale, growing in complexity with each emerging processor design. It needs to span a large distance covering the entire processor core, networking an increasing number of devices.

In order to improve the accuracy of the thermal monitoring systems, intensive research has been devoted to developing precise thermal sensing devices [4, 5]. However, the relationship between the supporting infrastructure and monitoring accuracy is not well understood. In this paper, we address the impact of the underlying infrastructure on the accuracy of the thermal monitoring system. We demonstrate that different methodologies for constructing this infrastructure can lead to drastically different accuracy. Therefore, we argue that the design of this infrastructure is at least as important for the overall accuracy and quality of the thermal monitoring as the individual sensing devices contained within this network.

We propose a novel optimization technique for this infrastructure. Our target thermal monitoring systems are those based on the remote sensing method, which is widely adopted in commercial processor designs. Particularly, we address the network delivering bias currents to the sensing devices. We developed a systematic optimization framework, where we formulate the network optimization problem as a bounded degree Steiner minimal tree problem on a metric graph.

We propose an ILP formulation for this problem.

Process variation is a major concern in ensuring the robustness and accuracy of the bias current network, particularly, since it contains analog components. We performed Monte Carlo simulations to validate the resulting routing infrastructure and quantify the impact of process and temperature variation on the accuracy of the monitoring system. The results show that the monitoring accuracy can be as high as 0.6°C under considerable amount of process and temperature variation. To evaluate the benefit obtained by introducing the on-chip current mirrors, we solve the ILP formulation of the problem and compare our solutions with a network with dedicated current sources for each thermal diode. We show that the reduction in the total wire length of the current distribution network can be as much as 42.74% and 27.65% on average.

The remainder of the paper is organized as follows. Section II gives an overview of the related work. In Section III, the interplay between the measurement infrastructure and the accuracy of the thermal monitoring system is analyzed, followed by the in-depth discussion of the routing network structure evaluated in this work. We introduce the optimization of the routing infrastructure in Section V. Experimental results are presented in Section VI.

II. RELATED WORK

High-performance microprocessors contain thermal monitoring modules in order to prevent the systems from entering severe thermal conditions. Dorsey et al. [3] described the thermal monitoring system for the AMD quad-core Opteron™ processor, where each core contains a number of remote temperature sensors scattered across the core and the sensor readings are routed to a central thermal evaluation unit. Duarte et al. discussed thermal sensing techniques used in an Intel® Pentium® 4 processor [6]. Both local and remote sensors are employed in the processor. Several sources of inaccuracy of the monitoring system are discussed. However, the implications of the temperature-dependent series resistance of interconnects on temperature measurement accuracy has not been addressed. In this work, we demonstrate that this can be a significant cause of inaccuracy. Furthermore, we perform a detailed analysis to assess its effect on different measurement methods including the 4-wire method. While there are discrete-component-based temperature sensors adopting the 4-wire principle [8], our contribution lies in the detailed evaluation of this method for on-chip temperature sensing in presence of process variations and thermal gradients. We also present a novel routing algorithm to generate a current distribution network for 4-wire measurement with minimal wirelength.

III. EXISTING APPROACHES AND CHALLENGES

Diodes are commonly used as thermal sensing devices. When biased by a forward current I_C , the forward bias voltage V_F of the diode depends reasonably linearly on the *absolute junction temperature* T_d in the proximity of the diode

$$T_d = \alpha V_F + \tau \quad (1)$$

where α is a constant of about -2.41mV/°C (the exact value depends on the saturation current of the diode and the forward current I_C) [9]; and τ is the interception of the linear function with axis $V_F = 0$.

In order to obtain the accurate temperature value, V_F should be precisely measured. A common method adopted by industrial designs




Figure 1. 2-wire voltage measurement.




Figure 2. The measured temperature as a function of ΔT_{avg} .




Figure 3. 4-wire (Kelvin) voltage measurement.

[6, 7] is the 2-wire measurement, which is depicted in Figure 1. Notice that V_F is measured and processed a certain distance away from the diode. The wire connecting the diode and the voltage measurement module is associated with *series resistance* R_s , which is a function of the temperature profiles along the length of the wire:

$$R_s = R_{s0} \cdot \left(1 + \beta \cdot \frac{\int_L (T(x,y) - T_0) dl}{\int_L dl} \right) \quad (2)$$

R_{s0} is the resistivity of the interconnect at room temperature T_0 (assumed to be 25 °C), β is the temperature coefficient of resistance and the integration is performed along the wire. In Equation (2), the ratio of the two integration terms can be interpreted as the average variation of temperature with respect to the nominal temperature along the interconnect, denoted by ΔT_{avg} .

The junction temperature of the diode is related to the measured voltage difference by V_{meas} and R_s by

$$T_d = \alpha (V_{meas} - I_C R_s) + \tau \quad (3)$$

Equations (2) and (3) indicate that without sufficient knowledge about the temperature distribution along the wire, the actual junction temperature cannot be precisely determined. The thermal diodes and the measurement circuitry can be placed very far away from each other, such that the interconnect between them would have to cover significant distances. These lines may need to cross an entire processor core [3]. It is certain that the thermal profile across such distances on a high-performance processor will vary significantly. Furthermore, a detailed map of this profile will not be available. Hence, the measurement error introduced by ΔT_{avg} *cannot be compensated by any kind of calibration*.

We performed HSPICE simulations to determine the relationship between the error in temperature measurement and ΔT_{avg} for the 65nm technology with copper interconnects. We set the width, and length of the interconnect between the diode and voltage measurement point to be 180nm and 3mm, respectively. The current value I_C is 10μA. In this experiment, we fix the temperature of the diode at 35°C and vary the temperature profile along the interconnect. Figure 2 plots the relationship between the measured temperature T_{meas} and ΔT_{avg} using a solid line. The measured temperature value is given by Equation (3). The fixed diode temperature (35°C), is also depicted as a reference using a dashed line. As ΔT_{avg} increases, we observe that there exists a significant disagreement between the actual and measured temperature. For instance, when ΔT_{avg} is 25°C, the measurement error can reach 5°C. In modern microprocessors, when running different applications, the die temperature can be dramatically different. Depending on the thermal throttling threshold, critical functional units such as integer/floating ALU, integer/floating register file, instruction queue, etc. can reach up to 90 °C. Temperatures of other functional units can be over 60 °C[10]. Hence, it is not unusual for ΔT_{avg} to be as large as 25°C.

IV. INFRASTRUCTURE FOR THERMAL SENSING

The high level topology of the sensing infrastructure supporting the alternative 4-wire measurement is depicted in Figure 3. Using the 4-wire measurement, the measured voltage V_{meas} is equal to the forward bias voltage of the diode. According to Equation (1),

$$T_d = \alpha V_{meas} + \tau \quad (4)$$

Unlike Equation (3), Equation (4) does not involve R_s . Therefore, the serial resistances of the wires and their thermal dependency do not impact the accuracy of the measurement. Figure 4(a) depicts the high level structure of the measurement network. The bias currents for thermal diodes are distributed through a network of current mirrors. Notice that the resulting current distribution network has a tree topology. In the rest of the paper, we will use the term *current mirror tree* as a synonym of this current distribution network structure. Figure 4(b) illustrates structure of a cascode current mirror built using PMOS transistors. The cascode current mirror structure has high output resistance, which eliminates the coupling between the current mirrors and the thermal diodes.

The 4-wire measurement method eliminates dependencies of thermal measurement accuracy on the wiring of the measurement network. However, this design is more resource-demanding, since the entire routing infrastructure contains the voltage measurement network in addition to the bias current distribution network. Furthermore, we note that the current mirrors could be sources of inaccuracy to the monitoring system. Firstly, due to process variation, the *matching ratio* of a current mirror (output current over input current) might not be exactly unity. As a result, the actual bias current of a thermal diode could deviate from the current generated by the current source. However, the impact of process variation on the system accuracy is *static*. In other words, the matching ratio of each individual current mirror is a fixed value after manufacturing. *Therefore, the impact of process variation on sensor accuracy can be mitigated by sensor calibration*.

Secondly, spatial and temporal temperature variation also affects the matching ratio of a current mirror. We performed Monte Carlo simulations to quantify the impact of process and temperature variations to the accuracy of the system. We observed that our monitoring system exhibits high level of robustness to the variations – the average accuracy of the system can be as high as 0.6 °C. Details of this analysis are presented in Section VI.B.

On the other hand, the interconnects for voltage measurement, would be directly routed from the thermal diode to the voltage measurement circuits as shown in Figure 4(a). Notice that crosstalk induced noise can be a source of inaccuracy for voltage measurement. However, this effect can be minimized by adding adequate shielding to the lines as it is currently practiced in industrial designs [6]. Furthermore, the thermal time constant is in the order of millisecond, which is several orders of magnitude larger than the clock cycle time (in the order of nanosecond). Thus, we can sample the voltage in multiple clock cycles and use the average voltage for temperature calculation. This would effectively filter out the inaccuracy imposed by the coupling between interconnects.

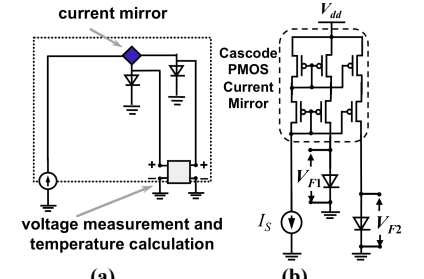


Figure 4. The bias current distribution network (a) a high-level view (b) transistor-level schematic

A. Extended Structure Accommodating more Thermal Diodes

The current mirror tree can be easily generalized to accommodate larger number of diodes. Figure 5 provides the transistor level schematic of a current distribution network, which can support up to

four diodes. Notice that we use alternating NMOS and PMOS current mirrors in the tree. This is necessary in order to create current flows in the appropriate directions. The input/output currents of the NMOS current mirror flow inward, while the input/output currents of the PMOS current mirror flow outward.

The structure of the current mirror is quite simple. Available whitespace on the layout could be utilized to embed them into the chip. The area of a current mirror is comparable to that of a repeater. Therefore, like the repeaters, the current mirrors can be inserted into the whitespaces of the chips [10]. Hence, the insertion of current mirrors can be performed in post-layout stage. Therefore, the necessary modification made to the design flow is minimal. Finally, in the above discussion, for simplicity, we assume a current mirror is able to map one input current to two identical currents. In fact, the current mirror can be extended to map one input current to more than two output currents by simply replicating the transistors. However, since the current mirrors are inserted into the whitespace whose capacity is limited, the number of output currents of a current mirror cannot be made arbitrarily large.




Figure 5. The current distribution network supporting four diodes.

V. OPTIMIZATION OF THE ROUTING NETWORK

The routing network described above supports higher precision compared to the routing network used by the 2-wire measurement. However, it has higher interconnect overhead. The problem of allocation and placement of the current mirrors within the network to obtain the minimum wirelength remains to be addressed. In this section, we will provide a systematic treatment of this problem.

A. Problem Formulation

As mentioned earlier, the current distribution network has a tree topology. More precisely, it is a Steiner tree with the diodes being the leaf nodes and the current mirrors being the Steiner points. Steiner tree is a well-known structure utilized in VLSI physical design, specifically for signal routing. The total length of the Steiner tree directly corresponds to the overhead of the routing network. Therefore, we are targeting on minimizing the total length of the network.

We can formulate the problem as a Steiner minimal tree problem on graphs. For a given chip, a weighted graph can be constructed in the following way: each diode/whitespace is represented by a vertex. The weight of an edge between two vertices equals to the Manhattan distance (assuming rectilinear routing) between the corresponding diodes/whitespaces.

The Steiner minimal tree problem on graphs is a well-studied problem [11, 12]. However, our problem is distinct from other variants in the literature, because in our problem, the Steiner points cannot have arbitrarily large degrees. As mentioned earlier, the capacity of each whitespace is limited. This means, the number of transistors, thereby the number of output currents of a current mirror occupying each whitespace cannot be made arbitrarily large. Therefore, each possible Steiner point is associated with an upper bound on the maximum degree on the Steiner tree. We mentioned above that we require the diodes to be the leaf nodes. This constraint is equivalent to requiring the maximum degree of the diode nodes to be one. Hence, in our problem, both Steiner points and leaf nodes, are associated with an upper bound on their maximum degrees on the Steiner tree. The problem can be defined formally as follows:

Definition 1. A *metric graph* is a positively weighted graph where the edge weights satisfy the triangle inequality.

Definition 2. The *degree* of a vertex v on graph G , denoted by $\deg_G(v)$, is the number of the edges adjacent to v on G .

Problem 1 (Bounded Degree Steiner Minimal Tree, BDSMT). Given a metric graph G , we denote its vertex set and edge set by $V(G)$ and $E(G)$, respectively. Each vertex $v \in V(G)$ is associated with a positive integer b_v called *degree bound*. A vertex set $U \subseteq V(G)$ is called the *terminal vertex set*. The bounded degree Steiner minimal tree problem seeks a minimum weighted subgraph T of G connecting all the vertices in U such that for each vertex v of T , $\deg_T(v) \leq b_v$.

It can be proven that Problem BDSMT is computationally intractable. We omit the proof due to space constraints.

B. ILP Formulation for Problem BDSMT

We propose an Integer Linear Programming (ILP) formulation for Problem 1 based on the existing ILP technique for the Steiner minimal tree problem on graphs [12]:

$$\text{minimize } \sum_{e \in E} w_e \cdot x_e, \text{ subject to} \quad (5)$$

$$\sum_{e \in \delta^-(S)} x_e \geq 1, \forall S \subset V(G), S \cap U \neq \emptyset, (V(G) \setminus S) \cap U \neq \emptyset \quad (6)$$

$$\sum_{e \in NE(v)} x_e \leq b_v, \quad \forall v \in V \quad (7)$$

$$x_e \in \{0, 1\}, \quad \forall e \in E \quad (8)$$

In the above formulation, w_e denotes the weight of edge e , and x_e is essentially an indicator of whether an edge should be included in the solution. Inequality (6) guarantees that the sub-graph consisting of edges whose $x_e = 1$ are connected, where $\delta^-(S)$ denotes the set of edges entering S . Inequality (7) bounds the degree for each vertex, where $NE(v)$ is the set of the neighboring edges to v .

VI. EXPERIMENTAL RESULTS

We performed Monte Carlo simulations to validate the functionality of the proposed infrastructure under process variation and different thermal gradients across the chip. We also experimented with the ILP formulation proposed for optimizing the wirelength of the current distribution network.

A. Validation Setup for the Routing Infrastructure

In our experiment, a total number of 1000 chips were examined subject to process and temperature variations. The chip simulations were carried out using HSPICE based on the IBM 10SF 65nm technology models [13]. In the simulation, each chip is assumed to be a 2mmx2mm square. We further divided the square into four 1mmx1mm quadrants and allocated one thermal diode at the center of each quadrant. The current distribution network shown in Figure 5 was used to provide the bias current for the diodes. The length of the global interconnects between the diodes and the current mirrors were in the range of 0.5mm to 2mm. For these wires, we have accounted for the dependence of their serial resistances on temperature in our HSPICE model. The “thermal diode” used in our simulation is actually an NMOS transistor with its drain and base connected, and gate and source shorted. To minimize the coupling between the current mirrors and the thermal diodes, we have properly sized the MOS transistors in the current mirrors and the diodes. The transistors in the current mirrors have small width/length ratio in order to increase the output resistance of the current mirror; whereas the transistors which are used as thermal diodes have large width/length ratio to reduce the input resistance of the diodes. In our experiment, the width/length ratios of the transistors in the current mirrors and the thermal diodes were chosen to be 1:4 and 80:1, respectively.

The simulation for each of the 1000 chips consists of three steps:

Parameter Perturbation: In the first step, the length and width of the transistors were perturbed, where spatial correlation was

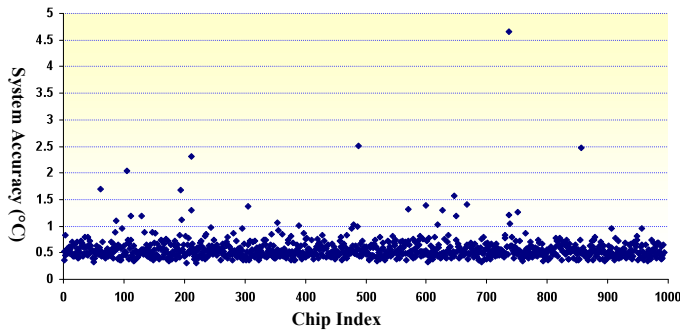


Figure 6. Monte Carlo simulations for the 3-point calibration scheme.

TABLE I. STATISTICS OF SYSTEM ACCURACY

	Avg Accuracy (°C)	Best Accuracy (°C)	Worst Accuracy (°C)	Stdev Accuracy (°C)
2-Point	1.082	0.691	4.976	0.397
3-Point	0.566	0.299	4.663	0.235

TABLE II. ROUTING WIRELENGTH RESULTS

Benchmarks		CMT	DCS	
Name	#diodes	opt-wl (mm)	wl (mm)	Impr%
Rbs1	2	12.25	13.59	9.86%
Rbs2	4	21.24	30.49	30.34%
Rbs3	8	40.43	70.61	42.74%

accounted for. This perturbation assumes Gaussian distribution. The standard deviation of the distribution was set to 15%.

Sensor Calibration: In the second step, the thermal diodes were calibrated. Both 2-point calibration and 3-point calibration were assessed. In the 2-point calibration process, we set the chip at uniform temperature levels 35°C and 115°C and perform HSPICE simulation. The forward bias voltages of each diode at these two temperature levels were measured. Linear fittings for the voltage-temperature curve of each diode were determined based on these two measurement points. In the 3-point calibration process, the chip was set under uniform temperature levels 35°C, 75°C, and 115°C. Piecewise linear fitting was used for temperature estimation. Since the relationship between the forward bias voltage of a diode and its junction temperature is actually a non-linear function, the 3-point calibration scheme is expected to have higher accuracy.

Monitoring Accuracy Evaluation: In the third step, the accuracy of the thermal monitoring system was evaluated using 10 randomly generated chip temperature profiles. The lower bound and upper bound of the temperature are 35°C and 115°C, respectively. For each temperature profile, the temperature in the proximity of each diode was estimated based on its forward bias voltage and the linear fitting (for 2-point calibration) and piecewise-linear fitting (for 3-point calibration) determined in the previous step. The *measurement error* of each diode is defined as the absolute value of the difference between the projected temperature and the actual temperature in the proximity of the diode. Further, we define the *system accuracy* of the monitoring system implemented on a chip as the maximum measurement error among the 4 diodes across the 10 temperature profiles.

The average/best/worst accuracy and the standard deviation of the accuracy are calculated for the batch of chips. Here the *average accuracy* is defined as the average of the accuracy across the 1000 chips. The best/worst accuracy and the standard deviation of the accuracy are defined similarly.

B. Validation Results

Figure 6 provides system accuracy for the 1000 chips for the 3-point calibration scheme. Table I summarizes the statistics of the system accuracy for the two calibration schemes. We observe that

our proposed infrastructure supports accurate thermal monitoring. For the 2-point calibration scheme, the average accuracy is around 1°C. In the 3-point calibration case, the average accuracy can be less than 0.6°C. The second finding is that our proposed infrastructure is immune to process and temperature variations. For both calibration schemes, the standard deviation of the system accuracy is small, which means the monitoring system on most of the chips can achieve the accuracy of 1°C or less.

Nonetheless, there are a few chips whose system accuracy is not as good (larger than 3°C). This is because for these chips, the perturbed channel lengths of some transistors were too small (less than 25nm). These transistors exhibit high order of non-linearity affecting the accuracy of the monitoring system.

C. Wirelength Optimization

We generated several benchmarks to evaluate the benefit of optimizing the routing of the current mirror tree. These benchmarks contain different number of sensors which are pre-placed on a microprocessor floorplan (Alpha EV6). During this placement, the goal was set as to minimize the measurement error for thermal profiles under typical workloads with a fixed number of sensors. Using this initial sensor location information we generated a bias current network for different number of sensors. The number of thermal diodes varies between 2 to 8.

Table II provides the experimental results. The column “#diodes” represents the number of diodes for each benchmark. “DCS” stands for *Dedicated Current Sources* for each diode, where each diode is biased individually. “CMT” denotes the network enhanced with the *Current Mirror Tree*. Column “wl” under “DCS” gives the total wirelength of the current distribution network for each benchmark (the voltage measurement network is not included in the calculation, since it is same in either design). We solve the ILP formulations of the BDSMT instances for each benchmark using a commercial ILP solver, CPLEX Ver10.1. The total wire lengths of resulting current distribution networks are listed in column “opt-wl”. The relative saving of wire length can be as much as 42.74%, and is 27.65% on average.

REFERENCES

- [1] Rotem, E., et al. *Analysis of Thermal Monitor Features of the Intel Pentium M Processor*. in Workshop on Temperature Aware Computer Systems. 2004.
- [2] Poirier, C., et al. *Power and Temperature Control on a 90nm Itanium-Family Processor*. in Intl. Solid-State Circuits Conf. Feb. 2005.
- [3] Dorsey, J., et al. *An Integrated Quad-Core Opteron™ Processor*. in Intl. Solid-State Circuit Conf. Feb. 2007.
- [4] Chen, C., et al., *A Time-to-Digital-Converter-Based CMOS Smart Temperature Sensor*. IEEE J. of Solid-State Circuits, 2005. 40(8): p. 1642-1648.
- [5] Chen, Q., et al. *A CMOS Thermal Sensor and Its Application in Temperature Adaptive Design*. in Intl. Symp. on Quality Electronic Design. 2006.
- [6] Duarte, D., et al. *Advanced Thermal Sensing Circuit and Test Techniques Used in a High Performance 65nm Processor*. in Intl. Symp. on Low Power Electronics and Design. Aug. 2007.
- [7] Pertijis, M. et al. *Transistor Temperature Measurement for Calibration of Integrated Temperature Sensors*. in Instrumentation and Measurement Technology Conference. May 2002.
- [8] *Improving The Accuracy of Temperature Measurements*. <http://www.picotech.com/applications/temperature.html>.
- [9] Ocaya, R., *An Experiment to Profile the Voltage, Current and Temperature Behaviour of a P-N Diode*. European Journal of Physics, 2006. 27: p. 625-633.
- [10] Chen, S., et al. *Floorplanning with Consideration of White Space Resource Distribution for Repeater Planning*. in Intl. Symp. on Quality Electronic Design. Mar. 2005.
- [11] *Steiner Trees in Industry*. 2001, Kluwer Academic Publishers. p. 235-279.
- [12] Lin, G. et al. *On the Terminal Steiner Tree Problem*. Information Processing Lett., 2002. 84: p. 103-107.
- [13] *IBM 10SF CMOS Process* in <http://www.mosis.com/ibm/10sf/>.