

Temperature-to-Power Mapping

Zhenyu Qi[†], Brett H. Meyer[‡], Wei Huang[§], Robert J. Ribando^{*}, Kevin Skadron[‡], Mircea R. Stan[†]

[†] Dept. of ECE, [‡] Dept. of CS, ^{*}Dept. of Mechanical and Aerospace Engineering, University of Virginia, Charlottesville, VA 22904

[§] IBM Austin Research Lab, Austin, TX 78758, USA, ^{†‡*}{jerry, bhm, rjr, ks7h, mircea}@virginia.edu, [§]huangwe@us.ibm.com

Abstract—Accurate power maps are useful for power model validation, process variation characterization, leakage estimation, and power optimization, but are hard to measure directly. Deriving power maps from measured thermal maps is the inverse problem of the power-to-temperature mapping, extensively studied through thermal simulation. Until recently this inverse heat conduction problem has received little attention in the microarchitecture research community. This paper first identifies the source of difficulties for the problem. The inverse mapping is then performed by applying constraints from microarchitecture-level observations. The inherent large sensitivity of the resultant power map is minimized through thermal map-filtering and constrained least-squares optimization. Choices of filter parameters and optimization constraints are investigated and their effects are evaluated. Furthermore, the paper highlights the differences between the grid and block modeling in the inverse mapping which were often ignored by previous schemes. The proposed methods reduce the mapping error by more than 10× compared to unoptimized solutions. To our best knowledge this is the first work to quantitatively evaluate and minimize the noise effect in the temperature to power mapping problem at the microarchitecture level for both grid and block mode, and for the steady and transient case.

I. MOTIVATION FOR POST-SILICON POWER MAPS

The “power wall” has become a critical performance limiter for integrated circuit design. Excessive power consumption leads to short battery life, higher utility costs, large currents in interconnect, and elevated temperatures. Moreover, the very power models which are intended to enable power-aware analysis and optimization are often inaccurate—in part because of changes in leakage power due to parameter variation—and need to be validated against silicon measurement. If they could be reliably derived, accurate *power maps* (the temporal and spatial power distribution) would provide an avenue not only to perform post-silicon power model validation, but also to characterize process variation over large regions, complementary to timing based methods like maximum frequency or critical path delay monitoring [1]. Per unit power consumption could also be used to inform power/energy-aware task scheduling or reconfiguration, and to capture long term wearout mechanisms and proactively eliminate potential reliability and thermal hazards.

Unfortunately, direct fine grain power measurement is expensive and difficult, and event-counter based power proxies cannot capture workload induced power variation [2]. However, accurate and high resolution *thermal maps* (temporal and spatial temperature distribution) can be measured using infra-red (IR) cameras or thermal sensors. While the *direct heat conduction problem* (DHCP) of solving thermal maps from power maps is essential for thermal simulation [3], [4], the *inverse heat conduction problem* (IHCP) of solving power maps from thermal maps is a possible approach to obtaining power maps without the need for actual power measurements. The importance of the microarchitecture level IHCP has only been recently recognized [5], [6], [7], [8]. While the power-to-temperature mapping is a pre-silicon design step, the temperature-to-power mapping is a “closing the loop” post-silicon task, deriving the very power maps needed to

characterize process variation, validate power models, and support new techniques to manage runtime power dissipation on a per-chip basis.

Though IHCP is “simply” the inverse of the well-studied DHCP, it is also known as a challenging problem [9], mainly due to the large sensitivity of power maps to thermal measurement noise (or uncertainty). Previous work on microarchitectural IHCP (named ‘MIHCP’ henceforth) [5], [6], [7] proposed algorithms without evaluating the impact of input temperature noise, which can come both from thermal measurement or from averaging and congregating. Fortunately, there are a number of characteristics of the MIHCP that make it more tractable than the general IHCP.

In this paper, we present a novel methodology for solving the steady-state temperature-to-power mapping problem in integrated circuits in the presence of thermal map noise. Our approach applies constraints and filters to the IHCP problem based on simple observations from real silicon systems. For example, all power dissipators must be in the silicon layers of a packaged chip, and spatial temperature must be continuous. This transforms the direct inverse problem into an over-constrained optimization problem, and, as a result, we are able to derive accurate power maps in the presence of thermal map noise, reducing error by up to 10× compared to unoptimized solutions. Moreover, unlike prior work, our methodology is applicable to both common types of compact thermal models prevalent in the literature, the block model and the grid model.

The IHCP is formulated in Section II and its challenges discussed in Section III. Section IV present unoptimized and optimized solution for the steady state. The transient case is analyzed in Section V. Section VI comments on related work and Section VII concludes the paper.

II. THERMAL CONDUCTION IN MICROARCHITECTURE

In microarchitecture studies, compact thermal modeling is extensively adopted [3], [4], [5] for easy modeling and fast computation compared to distributed models that have to be solved by finite difference or finite element methods. The steady-state heat conduction is modeled as

$$A_R \cdot P = T, \text{ or } \begin{pmatrix} a_{11} & a_{12} & \dots & a_{1n} \\ a_{21} & a_{22} & \dots & a_{2n} \\ \dots & \dots & \dots & \dots \\ a_{n1} & a_{n2} & \dots & a_{nn} \end{pmatrix} \begin{pmatrix} p_1 \\ p_2 \\ \dots \\ p_n \end{pmatrix} = \begin{pmatrix} t_1 \\ t_2 \\ \dots \\ t_n \end{pmatrix} \quad (1)$$

P and T denote power and thermal maps respectively. A_R is the thermal resistance matrix. The name comes from the well known electrical-thermal analogy in compact modeling, where heat conduction is modeled with lumped RC (resistors and capacitors) circuits, and temperature and power are analogous to voltage and current, respectively. Capacitors are missing in Eq.1 since they have no effect in steady state (the transient case will be discussed later). A_R can be obtained either through measurement [5], [8] or modeling [3].

In Eq.1, p_i and t_i ($1 \leq i \leq n$) are the power and temperature for a geometric node i in the floorplan. A node represents

a functional block or a geometric grid cell. Eq.1 states that temperature at node i (t_i) is influenced by power dissipated by nodes (p_1, p_2, \dots, p_n) with weighting factors a_{i1}, \dots, a_{in} . The direct problem, DHCP, tries to find T from P . The inverse problem, IHCP, tries to solve P from T . The IHCP:

$$P = A_R^{-1}T = A_C T \quad (2)$$

is mathematically straightforward but turns out to be challenging in the presence of thermal noise.

III. PROPERTIES AND CHALLENGES OF MIHCP

The temperature-to-power mapping is known as an ‘ill-posed’ problem in that the solution is very sensitive to input noise [9]. This high sensitivity at the microarchitecture level has not been discussed in previous work [5], [6], [7]. Appropriate optimization can only be applied once the source of the sensitivity is understood. This section dismisses two plausible reasons and identifies the source of the large sensitivity in the steady state MIHCP described by Eq.2.

The first common difficulty in solving a linear system like Eq.2 is the large matrix size. Iterative methods are often used to avoid intractable computation time, but could have convergence problems that result in solution instability. However, extremely large thermal conduction/impedance matrices are most likely not needed in MIHCP. Recent work [10] shows that when power source sizes are 1.5 times larger than the silicon die thickness, the die can be modeled as a homogeneous 2-D plane with all the metal layers and detailed circuitry ignored. Furthermore, for state-of-the-art microprocessors with realistic packaging, power map granularity of around 400 microns is sufficient for 1°C resolution in thermal maps. For many applications mentioned in Section I even coarser power maps are sufficient.

Even with small sizes, taking the inverse or decomposition of an ‘ill-conditioned’ or singular matrix can lead to large sensitivity or non-uniqueness in solution due to numerical instability or singularity. The numerical stability of a matrix can be characterized by the *condition number* (the ratio of the largest to the smallest singular value). Large condition numbers mean small input errors can cause large output errors. However, a matrix has the same the condition number as its inverse. Thus the thermal conduction and resistance matrix A_C and A_R would cause the same instability in DHCP and IHCP. Since the large sensitivity to input noise does not exist in DHCP, it cannot be solely explained by numerical instability in the inverse problem.

The following example demonstrates the large sensitivity in the MIHCP even without large matrix sizes and numerical instability problems. The simulation setup is explained in Section IV-A. We first apply $P=(26 \ 97 \ 150)$ (Watts) to the three units respectively in the floorplan in Fig.1(a), and solve the steady state temperature to be $T=(41.7 \ 46.7 \ 65.7)$ with Eq.1 (without loss of generality we set the ambient temperature to 0 (Kelvin) for all simulations in this paper, so only the temperature raise due to on-chip power dissipation is shown). Slightly perturbing power numbers to $P=(25.8 \ 94.3 \ 150.3)$ and solving again for T yields a change of less than 2%. Next we start from T and solve for P , with the intent of arriving at, or close to, the original power map. Without thermal noise the obtained power map matches the original values perfectly with Eq.2. However, when the thermal map is slightly perturbed to $T=(41.4, 45.4 \ 65.8)$, the power map becomes $P=(4.8 \ 60.7 \ 103.2)$, with a mean error of more than 50% and maximum error of more than 80%!

What is the reason for the large sensitivity? A close examination of the above and other examples points to the relatively large elements with opposite signs of elements in A_C . The intuition is that power can change abruptly both




Fig. 1: Two floorplans in HotSpot download package.

spatially and temporally, but temperature cannot. Trying to reconstruct high frequency components in power maps from low frequency thermal maps renders the problem extremely sensitive to input thermal noise even without large matrix sizes or condition numbers. The steady-state analysis above only demonstrates sensitivity due to spatial high frequency in power maps. The transient analysis in Section V further addresses temporal high-frequency induced sensitivity.

IV. SOLVING MIHCP IN STEADY-STATE

Our goal is to solve the power map with minimum error when thermal map uncertainty or noise are present. While IHCP is very sensitive to thermal-map noise, there are a number of opportunities to use additional information—constraints—specific to the inverse problem in a packaged chip to manage the effect of that noise on the derived power map. These constraints all proceed naturally from observations about packaged chips, and include, for example, the fact that all power is dissipated only in the silicon layers. We have performed a variety of simulations to quantify the effect of these constraints in isolation and collectively, all in comparison with a direct solution of the MIHCP.

A. Simulation Setup

In our simulations, we first use HotSpot [3] to derive the thermal conduction matrix A_C , given a floorplan and packaging details. We then use Matlab (2009b) to (1) derive the temperature map T using the power map P , (2) permute T with Gaussian distributed noise, and (3) derive the observed power map P' . All of our constraints and optimizations are implemented in Matlab.

We elected to use a compact thermal model to generate A_C for several important reasons. Previous work on inverse mapping obtained thermal maps from either IR cameras or thermal sensors [5], [6], [8], [11]. In either case, errors cannot be quantified without also knowing the input power map, which is challenging for aforementioned reasons. Employing a compact thermal model gives us the flexibility to evaluate the influence of variable thermal map noise on the derived power maps relative to the original power map. Of course, our optimizations can be applied to temperature measurements derived in any fashion.

A compact thermal model also gives us the flexibility to explore the effect on power map derivation of different approaches to thermal measurement, such as the choice of a *block mode* or *grid mode* thermal model. *Block mode* supports non-uniform node sizes in silicon, while *grid mode* imposes a uniform, fixed node size. The former is a better proxy for the nonuniformly measured thermal maps from sensors placed over critical components or hotspots, while the latter approximates the uniformly measured thermal maps from IR cameras or uniformly distributed thermal sensors. HotSpot can generate the appropriate A_C in either case.

B. Unoptimized MIHCP

We begin our simulations using the $22cm \times 22cm$ floorplan of a testchip developed at UCSC [12], containing 10×10 identical square blocks. We randomly assign power values

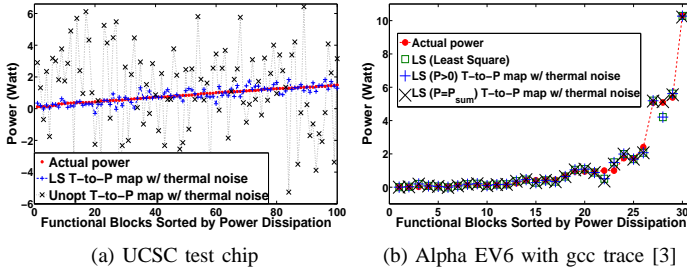


Fig. 2: Steady state inverse mapping. Thermal noise $\mu = 0, \sigma = 5\%$

varying from 0.1W to 1.5W to all 100 blocks and the steady state thermal map is obtained with Eq.1. Thermal noise is modeled with a normal distribution with mean (μ) set to 0, since constant offsets can be removed through calibration. Fig.2(a) illustrates the large error that results from Eq.2 (Unopt vs. Actual Power). The x and y axes denote the blocks and power values for each block respectively.

C. Initial Optimization Using Least Squares

Applying Eq.2 without any optimization yields large errors as shown in Section III and Fig.2(a). Our first observation is that none of the nodes from the TIM (thermal interface material), the heat spreader and the heat sink dissipate power. In reality temperature at these nodes may be obtained with a thermal testbed. We can then rewrite Eq.2 in a block fashion:

$$\begin{bmatrix} A_{R,si} & A_{R,12} \\ A_{R,21} & A_{R,pk} \end{bmatrix} \begin{bmatrix} P_{si} \\ P_{pk} \end{bmatrix} = \begin{bmatrix} T_{si} \\ T_{pk} \end{bmatrix} \quad (3)$$

where subscript si denotes nodes on the silicon die and pk denotes nodes associated with cooling and packaging that have zero power dissipation. P_{pk} contains only zeros and can be removed. least-squares (LS) optimization can be applied to the following over-constrained equation:

$$\begin{bmatrix} A_{R,si} \\ A_{R,21} \end{bmatrix} \begin{bmatrix} P_{si} \end{bmatrix} = \begin{bmatrix} T_{si} \\ T_{pk} \end{bmatrix} \quad (4)$$


Fig.2(a) illustrates the improvement from applying the least-squares (LS) optimization by Eq.4. Detailed data are reported later in Section IV-F.

Without thermal noise, the results from Eq.2 match the original power map perfectly. With thermal noise, the uncertainty gets amplified by the large entries in A_C leading to large power map errors. On the other hand, the result of LS is constrained in a subspace that contains the noise-free solution and the thermal noise simply perturbs that solution in the subspace.

D. Finding More Constraints for Optimization

Turning Eq.1 into an *over-constrained* problem by utilizing microarchitectural information significantly improves the MIHCP solution. Additional properties can be exploited to further improve the optimization. Constraints can be thought of as guiding the optimization away from undesired results.

We observe in Fig.2(a) that unconstrained LS results in some small but physically impossible negative power numbers. Furthermore, although power monitoring at the block level incurs design costs, the total chip power consumption is usually available at the power pins, and can serve as another constraint. We incorporate these constraints into the LS solver in Matlab, examine their effects. Fig.2(b) plots the mapping results for EV6 floorplan with 30 functional units shown in Fig.1(b). Noticeable mapping improvement is achieved on node 28 from the constrained LS. Detailed results are reported later in Section IV-F. The constraints not only help eliminate negative or non-physical values individually, but also improve the solution as a whole.



(a) Thermal noise $\mu = 0, \sigma = 5\%$

(b) Thermal noise $\mu = 0, \sigma = 25\%$

Fig. 3: Gaussian filter parameter study with different thermal noise. Optimum exists for σ . Odd h better than even numbers.

Lower and upper bounds of the power values for each block can also be added to the collection of constraints. However, it is not clear how conveniently or accurately these values could be obtained, and they are not included here. In general, the more constraints are applied in the optimization, the more likely that the results are close to actual values, although usually at a larger computational cost, which is discussed in Section IV-I.

E. Thermal Map Preprocessing

Power can change abruptly from block to block, but temperature has to be continuous and changes gradually over space. Therefore a low-pass filter can be applied to the thermal map to get rid of high frequency thermal noise before the inverse mapping. This has to be performed with knowledge of the floorplan since the 2-D spatial information is lost in the 1-D representation in Eq.1. In image processing, Gaussian filters are often used to ‘blur’ images to get rid of image details and noise [13]. The impulse response of a Gaussian filter is the standard Gaussian function. For a digital Gaussian filter H in Matlab, two parameters need to be specified, i.e., the filter size h (a square 2-D filter is assumed) and the standard deviation σ_f of the Gaussian function.

Fig.3 shows the filter effect on thermal map noise with varying filter parameters for the UCSC test chip. First, it is seen that the low pass filter is more effective with large thermal noise. For carefully chosen parameters, noise reduction is over 70% for $\sigma = 25\%$, and 20% for $\sigma = 5\%$. On the other hand, bad parameters can actually increase the error, especially when only small noise exists. For example filtering a noise-free thermal map would only introduce errors. Second, the optimal value for σ_f increases with the thermal noise σ , indicating that with large thermal noise more weight should be given to neighboring nodes for filtering. Third, odd values of h are better than even ones due to the symmetry in the filtering. Finally, optimal values of h increase with thermal noise for the same reason as σ_f : the error almost flattens out for large odd-numbered h . In fact the error increases slowly as h gets larger than 11 in both figures in Fig.3, since the chip is divided into a 10×10 grid and any weight on non-existent nodes is likely to only increase error. Therefore the optimum h should be set to the odd number closest to the grid size, in this case 11.

F. Steady State MIHCP Optimization Results

We first consider two extreme cases. Thermal map noise with $\sigma = 1.67\%$ (or $3\sigma = 5\%$, meaning with 99% confidence level the noise is within 5% of the true value) represents a highly accurate thermal map, which can be the case for some IR cameras [5], [6] or well designed thermal sensors [5], [11]. The other extreme with $\sigma = 25\%$ is also considered for demonstration and extrapolation purposes. Practical cases can fall anywhere in-between.

Five metrics are included in Table I where mapping results are reported for the UCSC test chip [12] with the

experimental setup described in Section IV-C. For unbiased evaluation, noise numbers are generated 1000 times and averaged data are reported. The 1st and 2nd columns show the maximal and average *absolute* error (in Watts). The 3rd column shows the average *relative* error. The relative error for large and small power nodes are evaluated in the last two columns. ‘Large power nodes’ are nodes that dissipate power more than half the maximal per block power, and ‘small power nodes’ are the rest. Note that the actual values of ‘avg rl’ usually decrease with a higher threshold in the ‘large power’ definition, simply due to the ‘relative’ nature. In practice, the power estimation for large power nodes is usually the most important. If the MIHCP is solved for power model validation or process variation characterization, a few power sources can be formed through scannable focused laser beams [5] or artificial benchmarks, which essentially create some ‘large power nodes’. Therefore, we focus more on reducing errors for these nodes.

algrm	max abs	avg abs	avg rel	avg rl	avg rs
un-opt	2.66/39.9	0.76/11.5	1.30/19.6	0.72/10.8	2.08/31.2
	0.40/5.95	0.057/0.86	0.13/1.93	0.072/1.08	0.28/4.25
LS 0	0.29/4.28	0.08/1.22	0.14/2.09	0.076/1.15	0.22/3.34
	0.044/0.65	0.064/0.10	0.014/0.22	0.0078/0.12	0.032/0.47
LS 1	0.28/3.40	0.081/0.81	0.14/1.19	0.076/0.85	0.22/1.65
	0.044/0.66	0.065/0.05	0.013/0.13	0.0078/0.068	0.029/0.30
LS 2	0.29/3.39	0.082/1.22	0.14/2.09	0.076/1.15	0.22/3.34
	0.043/0.66	0.066/0.10	0.014/0.22	0.0078/0.12	0.032/0.47
LS 3	0.29/3.39	0.081/0.81	0.14/1.19	0.076/0.85	0.22/1.65
	0.044/0.66	0.064/0.05	0.013/0.13	0.0078/0.068	0.029/0.30
LS 3F	0.28/2.35	0.081/0.67	0.14/1.10	0.076/0.65	0.22/1.70
	0.043/0.41	0.066/0.045	0.014/0.12	0.0077/0.058	0.029/0.27

TABLE I: Inverse mapping error by different algorithms for UCSC test chip. Thermal noise $\mu=0$, $\sigma=1.67\%/25\%$ (reported on left/right in each entry). ‘Un-opt’ = un-optimized inverse mapping. For LS, 0 = no constraint, 1 = constraint ‘ $P \geq 0$ ’. 2 = constraint ‘ $P = P_{sum}$ ’. 3 = both constraints. 3F = both constraints and filter. Mean/deviation are reported in upper/lower rows for each algorithm. *max/avg* denote maximal/average. *rl/rs* denote relative error for large/small power nodes. All values are obtained from 1000 random samples.

Compared to directly solving Eq.1, unconstrained LS alone (Eq.4) reduces errors by almost $10\times$. Constrained LS with non-negative power ($P \geq 0$) further leads to another 26% reduction for large power nodes in for very noisy thermal maps ($\sigma=25\%$). This is because the large solution sensitivity often leads to negative values in the solution when significant noise exists in input thermal maps, which can be seen in both examples in Fig.2. The $P \geq 0$ constraint not only removes the negative values but also improves the solution on the whole. For this example the total power dissipation constraint ($P = P_{sum}$) has little effect both by itself and when combined with $P \geq 0$, regardless of the thermal noise level. Finally, when filtering is applied before the constrained LS, the error is reduced by another 24% for large power nodes with $\sigma=25\%$ thermal noise. As discussed in Section IV-E, filtering is much more effective when significant thermal noise is present. Notice that filtering may reduce error for large power nodes at the expense of increased error for small power nodes, since the large error added on the peaks of thermal maps is likely to be smoothed out, but the smoothing effect can potentially alter small temperature values. Finally, the deviation values in Table I are also important since they indicate the confidence interval of mapping results.

Fig.4(a) shows how power map error changes with the noise magnitude in the thermal map for the UCSC test chip. Again we focus more on the large power nodes for reasons mentioned earlier, so the power map noise shown here is essentially the ‘avg rl’ in Table I. The unoptimized inverse




Fig. 4: Power map noise vs. thermal map noise in block mode inverse mapping. Unoptimized results omitted due to large errors.

mapping is not included since its error is more than $10\times$ larger. It can be seen that the power map noise increases almost linearly with thermal map noise. Constrained LS and spatial filtering reduce mapping error effectively across all thermal noise levels.

The same simulations are performed for the irregular Alpha 21364 (EV6) floorplan using the ‘gcc’ power trace mentioned in Section IV-D. The results are plotted Fig.4(b) shows much smaller power map error than Fig.4(a), due to the smaller number of total nodes and ‘large power nodes’. Notice that for this example we lowered the threshold of ‘large power’ from 1/2 to 1/3 of the maximal per-node power to include more nodes in the error characterization. The four units considered ‘large power nodes’ are L2 (level-2 cache), I-Cache (instruction cache), D-Cache (data cache) and IntExec (integer execution unit), which together dissipate 64.3% of the total steady state power. For thermal noise with $\sigma = 1.67$ the relative error for large power nodes and all nodes are 2.77% and 4.80% respectively.

We observe in both Fig.4(b) and Fig.2(b) that the $P = P_{sum}$ constraint now becomes more effective than $P \geq 0$, which is reversed from Fig.4(a). Intuitively, when power is dissipated from only a few nodes, as the case in Fig.2(b) or when laser beams [5] and artificial benchmarks are used, $P = P_{sum}$ pushes the optimization towards the true values. However, when the power dissipation is distributed more uniformly as in Fig.2(a), errors in the power map caused by the zero-mean thermal noise are more likely to cancel each other when summed together, rendering $P = P_{sum}$ less effective. Nonetheless $P \geq 0$ is very effective when large thermal map noise leads to negative values with unconstrained LS. Therefore it is always safe to apply both constraints if possible. Finally, notice that spatial filtering cannot be applied directly to the irregular floorplans like EV6.

G. MIHCP in Grid Mode

Having one power number for each functional block is spatially ‘biased’, since block areas vary significantly. For example, in Fig.1(b), L2 is extremely large compared to other blocks but is underrepresented by a single node in the block-mode thermal and power map in Eq.1. To obtain more spatially uniform information across the whole chip, and get insight into power dissipation within large blocks and capture local hotspots with large power density, the inverse temperature-to-power mapping can be carried out in grid mode. This section studies the MIHCP in grid mode with the irregular EV6 floorplan for different grid granularities.

The steady state temperature is first computed in block mode, and then in grid mode with 8×8 and 16×16 grid sizes. The ‘large power’ threshold is still set to 1/3, yielding 9 and 17 grid cells for the 8×8 and 16×16 cases respectively. None of these cells fall into the big L2 block now, since the power of a block is divided equally among the grid cells covering it, and the power density of L2 is low. Fig.5 plots grid mode




Fig. 5: Power map noise vs. thermal map noise in grid mode inverse mapping (block mode results plotted for reference).

mapping results. It can be seen that using grid mode itself does not improve mapping accuracy. However, the low pass filter can now be applied to the regular floorplan and its effectiveness is demonstrated again.

H. Power Map Conversion Between Block and Grid Modes


Power maps in both block and grid modes are useful. For example, power models or measured data from power proxies and performance counters can be validated against the former. Knowledge of process variation or the within-block power profile is better represented by the latter. For power-to-temperature mapping, thermal simulators like HotSpot often perform computation in grid mode for accuracy, and then convert results to block mode if it is desired. The conversion simply extracts the maximal or average grid temperature in the block as the block temperature.

However, special caution should be taken for the inverse mapping. For instance, with the gcc power trace and grid sizes of 8×8 , 16×16 and 32×32 , simply assuming equal power density and taking the average of overlapping grids to convert the grid mode power map to block mode incurs 26.51%, 17.20% and 9.76% error for D-Cache, the largest power dissipator, even though it covers more than 30 grid cells in the 32×32 case. This is again due to the discontinuity of power maps. For example, in the 32×32 case, even though the D-Cache fully covers 30 grid cells internally, it still has 19 boundary cells, 7 of them shared with the large L2 with power density 40 times smaller than the D-Cache. The conversion is less of a problem for thermal maps.

I. Analysis on Computational Complexity

Spatial filtering of a thermal map with n units takes a constant time $O(n^2)$ in the worst case. The computational complexity of our scheme is dominated by the LS optimization. In our current implementation, both the unconstrained and constrained least-squares problems are solved in Matlab. The exact computational complexity is difficult to quantify since Matlab uses different algorithms for problems of different scales and sparsity. Applying constraints incurs overhead since iterative algorithms like quadratic programming and preconditioned conjugate gradient (for medium and large scale problems respectively) are used. However in the worst case the computational complexity is bounded by $O(n^3)$, the cost of the direct inverse of a dense matrix.

Our simulations are performed on a laptop with Intel Core Duo CPU 2.0GHz, 2GB RAM. For the 16×16 grid mode for the EV6 floorplan which results in A_C of 1036×1036 , the unconstrained LS spends around 0.1s CPU time on solving Eq.4 once. Applying both constraints of $P \geq 0$ and $P = P_{sum}$ increases the cpu time to 11s if algorithms of medium scale are used. Enabling large scale algorithms reduces the cpu time drastically to 0.3s but increases the error on large power nodes by about 8%. All data in Table I are obtained through medium scale solvers to show what can be achieved. However, mapping accuracy can be traded for computation speed, which also depends on the implementation platform.



(a) Directly solved power trace of unit 1 w/o thermal noise. (b) Noise on $\frac{dT}{dt}$ is the main source for inverse mapping error.

Fig. 6: Transient inverse mapping without and with noise.

V. MIHCP IN THE TRANSIENT CASE

Dynamic power model validation or management might desire transient power maps, which could be obtained by three different methods: (1) solving the steady state MIHCP contiguously as in [5], [6]; (2) solving the steady state MIHCP for each architecture event individually and combining them with event counters; (3) directly solving the transient MIHCP from transient thermal maps. Methods 1-2 only require solving the steady state MIHCP which has been addressed in Section IV. While we propose method (2) as a better approach to (1), this section identifies the challenges and future research directions for the transient MIHCP, which has not been discussed extensively in prior studies.

Transient heat conduction can be modeled as

$$C \frac{dT}{dt} + A_C T = P, \quad (5)$$

where C is a diagonal matrix containing thermal capacitors. A_C is the same as in Eq.1. Eq.5 can be used directly for MIHCP, with $\frac{dT}{dt}$ approximated with the forward, backward or central difference methods. We choose the central difference method because of its smaller truncation error.

We start with the floorplan in Fig.1(a). Both units 2 and 3 dissipate 50 Watts power constantly, while unit 1 is has power spikes shown in Fig.6(a). An important observation is that even without thermal noise, the derived power trace does not perfectly match the original. The largest mismatch is more than 10% at the first peak and there is some overshoot at the end of the second spike. The inherent error is from the approximation of $\frac{dT}{dt}$ and does not exist in steady state.

In Fig.6(b), thermal noise of $\mu=0$, $\sigma=1.67\%$ is added. Similar to the steady state case, the noise amplification effect in the inverse mapping is obvious. Different from the steady state case, the noise actually has dual effects in Eq.5. It not only perturbs the temperature T , but also its derivative $\frac{dT}{dt}$. To evaluate their individual impact, we examine how the power map would change if we were able to remove the noise in $\frac{dT}{dt}$. In fact, simply applying the central difference method on the original T significantly improves the result in Fig.6(b), revealing that reducing the error in $\frac{dT}{dt}$ is the key to solving the transient MIHCP.

However, the large sensitivity to $\frac{dT}{dt}$ poses practical difficulty. Even a high frame rate IR camera can only achieve 125fps [6] (8ms/frame), while modern microprocessors run at multi-GHz (clock cycle less than 1ns). The simulation in Fig.6(b) is performed with thermal data of 0.33MHz, much faster than IR imaging. On the other hand, thermal sensors usually have an inherent tradeoff between accuracy and speed. A newly reported thermal sensor achieves temporal resolution of 1ms with 5°C maximum error with process and power supply variation [11]. While this combination of speed and accuracy is fast enough for thermal throttling, the above result shows that it may not be enough for transient temperature-to-power mapping. Therefore, transient power

maps could only be obtained from steady state power maps as discussed in the beginning of this section.

If reasonably accurate $\frac{dT}{dt}$ values are available, in principle all techniques in the steady state MIHCP could be applied. Moreover, filtering can be applied temporally since temperature at any node must change smoothly. By splitting the silicon and package nodes in Eq.5 and getting rid of package nodes, a LS problem can be formulated as:

$$\begin{bmatrix} A_{R,si} \\ A_{R,21} \end{bmatrix} \begin{bmatrix} P_{si} \end{bmatrix} = \begin{bmatrix} A_{R,si}C_{si} & 0 \\ 0 & A_{R,pk}C_{pk} \end{bmatrix} \begin{bmatrix} dT_{si} \\ dT_{pk} \end{bmatrix} + \begin{bmatrix} T_{si} \\ T_{pk} \end{bmatrix} \quad (6)$$

VI. RELATED WORK

For DHCP at the microarchitecture level, extensive research has been performed and tools developed, yet MIHCP has not caught much attention until very recently.

Parameter fitting based on the genetic algorithm (GA) is proposed for MIHCP in [6]. In this framework, mapping accuracy and granularity increase with parameter number, but meanwhile the search space and thus the complexity of GA increases exponentially. The GA based method targets MIHCP only in the block mode and would become intractable for the grid mode. Moreover, significant work on benchmarking and data collection are needed for GA to explore the search space. A maximum discrepancy of 75% is observed for total power consumption in a single frame.

The work in [7] proposes to derive power maps from thermal maps with image processing algorithms. Finite element simulations are still required and iterative solvers are needed. Treating power and thermal maps as images loses microarchitectural information on functional blocks and is valid only for grid mode. As pointed out in Section IV-H, conversion between the two modes can introduce large errors, but a very fine grid would significantly increase runtime. A maximum of more than 60% relative error for peak values in the derived power map is reported.

An adaptive resolution multigrid algorithm is proposed in [14] for solving IHCP for 3D devices. A number of power maps are obtained by varying device parameters and the one closest to the measured power is chosen. The algorithm is computationally intractable for MIHCP due to the large number of physical parameters and heat sources.

The SIMP methodology [5] from IBM captures thermal maps from IR cameras and solves power maps through LS with solutions constrained in certain ranges. It is not explained how the LS problem is formulated and what constraints are applied, both of which are described at length in this paper. More importantly we qualitatively evaluate the individual and combined effects of these constraints on power map noise minimization.

The concurrent work [8] published at the same time when this paper is being finalized also proposes thermal map filtering and optimization techniques for reducing error in the derived power maps. Besides the differences in the thermal map filtering and optimization schemes, the experimental setting in [8] makes it difficult to evaluate how output power map noise tracks input thermal map noise. On the contrary, using a thermal simulator as in this work provides more flexibility and better knowledge in evaluating the mapping optimization effectiveness with different noise types and levels. Moreover only a regular grid floorplan is evaluated in [8], while we further point out the importance of distinguishing grid and block mode mapping. In addition we also discuss the transient case IHCP and demonstrate its major challenge.

To our best knowledge this is the first work to quantitatively evaluate and minimize the noise effect in the temperature to power mapping problem at the microarchitecture

level for both grid and block mode, and for the steady and transient case.

VII. CONCLUSIONS AND FUTURE WORK

Fine-grain power maps at the microarchitecture level are useful for many purposes including validating power models, characterizing parameter variation, estimating leakage and maximizing power efficiency. The general temperature-to-power mapping is an ‘ill-posed’ problem due to the large sensitivity of solved power values to input thermal noise. This sensitivity arises from the large range of power values in the power map. By leveraging properties at the microarchitecture level we have shown that careful thermal map filtering and constrained least-squares optimization significantly reduce power map error when thermal map noise is present: compared to unoptimized mapping, our approach can reduce mapping errors by more than 10×. Our approach is applicable to both block and grid models. However, unlike for temperature maps which are spatially continuous, conversion between blocks and grids for power maps could introduce large errors. We further point out that the key in solving the transient MIHCP with input thermal noise is to address the error in $\frac{dT}{dt}$ rather than simply T .

VIII. ACKNOWLEDGMENTS

This work was supported in part by NSF grant CRI-0551630 and SRC grant 1607.001. We thank the anonymous reviewers for their helpful comments.

REFERENCES

- [1] A. Keshavarzi, J. W. Tschanz, S. Narendra, V. De, W. R. Daasch, K. Roy, M. Sachdev, and C. F. Hawkins, “Leakage and process variation effects in current testing on future CMOS circuits,” *IEEE Design and Test of Computers*, vol. 19, pp. 36–43, 2002.
- [2] M. Ware, K. Rajamani, M. Floyd, B. Brock, J. C. Rubio, F. Rawson, and J. B. Carter, “Architecting for power management: The POWER7 approach,” in *Proceedings of Int’l Symposium on High-Performance (HPCA)*, 2010.
- [3] K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan, and D. Tarjan, “Temperature-aware microarchitecture,” in *International Symposium on Computer Architecture*, 2003, pp. 2–13.
- [4] P. Li, L. T. Pileggi, M. Asheghi, and R. Chandra, “Efficient full-chip thermal modeling and analysis,” in *Proc. Int. Conf. on Computer Aided Design (ICCAD)*, 2004, pp. 319–326.
- [5] H. F. Hamann, A. Weger, J. A. Lacey, Z. Hu, P. Bose, E. Cohen, and J. Wakil, “Hotspot-limited microprocessors: Direct temperature and power distribution measurements,” *IEEE J. Solid-State Circuits*, vol. 1, no. 42, pp. 56–65, Dec 2007.
- [6] F. J. Mesa-Martinez, J. Nayfach-Battilana, and J. Renau, “Power model validation through thermal measurements,” *International Symposium on Computer Architecture*, vol. 35, no. 2, pp. 302–311, 2007.
- [7] X. Wang, A. Shakouri, S. Farsiu, and P. Milanfar, “Extraction of power dissipation profile in an IC chip from temperature map,” in *Annual Thermal Measurement, Modeling and Management Symposium*, 2007, pp. 51–56.
- [8] S. R. Ryan Cochran, Abdullah Nowroz, “Post-silicon power characterization using thermal infrared emissions,” in *Proc. Int. Symp. on Low Power Electronics and Design (ISLPED)*, 2010.
- [9] J. V. Beck, B. Blackwell, and C. R. S. Clair, *Inverse Heat Conduction: Ill-Posed Problems*. Wiley-Interscience, 1985.
- [10] K. Etezzam-Yazdani, H. F. Hamann, and M. Asheghi, “Impact of power granularity on chip thermal modeling,” in *Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems*, 2006.
- [11] E. Saneyoshi, K. Nose, M. Kajita, and M. Mizuno, “A 1.1V 35μm × 35μm thermal sensor with supply voltage sensitivity of 2°C/10%-supply for thermal management on the SX-9 supercomputer,” in *Symposium on VLSI Circuits*, 2006, pp. 9–10.
- [12] E. Ardestani, F. J. Mesa-Martinez, and J. Renau, “Cooling solutions for processor infrared thermography,” in *Annual Thermal Measurement, Modeling and Management Symposium*, 2010.
- [13] L. G. Shapiro and G. C. Stockman, *Computer Vision*. Prentice Hall, 2001.
- [14] P. E. Raad, P. L. Komarov, and M. G. Burzo, “Coupling surface temperature scanning and ultra-fast adaptive computing to thermally fully characterize complex three-dimensional electronic devices,” in *Annual Thermal Measurement, Modeling and Management Symposium*, 2006, pp. 204–209.