

Power-Aware Deployment and Control of Forced-Convection and Thermoelectric Coolers

Mohammad Javad Dousti and Massoud Pedram

Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA
{dousti, pedram}@usc.edu

ABSTRACT

Advances in the thermoelectric cooling technology have made it one of the promising solutions for spot cooling in VLSI circuits. Thermoelectric coolers (TECs) generate heat during their operation. This heat plus the heat generated in the circuit should be transferred to the ambient environment in order to avoid high die temperatures. This paper describes a hybrid cooling solution in which TECs are augmented with forced-convection coolers (fans). Precisely, an optimization framework called OFTEC is presented which finds the optimum TEC driving current and the fan speed to minimize the overall power consumption of the cooling system while maintaining safe die temperatures. Simulation results on a set of eight benchmarks show the benefits of the proposed approach. In particular, a baseline system without TECs but with a fan could meet the thermal constraint for only three of the benchmarks whereas the OFTEC solution satisfied thermal constraints for all benchmarks. In addition, OFTEC resulted in 5.4% less average power consumption for the aforesaid three benchmarks while lowering the maximum die temperature by an average of 3.7°C.

Categories and Subject Descriptors

B.7.1 [Integrated Circuits]: Types and Design Styles – *advanced technologies, VLSI (very large scale integration)*.

Keywords

Dynamic thermal management, thermoelectric coolers, forced-convection cooling, cooling, low-power design, leakage power.

1. INTRODUCTION

Decreasing the transistor feature size has led to increased power density in VLSI circuit substrates. High power density causes hot spots on the chip, which tends to accelerate the device and interconnect aging processes and may even cause permanent physical defects if the temperature of these hot spots exceeds a certain threshold. Additionally, increased die temperatures result in slower devices and higher leakage power dissipation. As a result, thermal issue is one of the main barriers to the successful continuation of Moore's law. The purpose of a thermal management system is to stop the temperature increase beyond a certain threshold, even if the required action is to power off the chip.

Different thermal management solutions have been proposed during the past decade. These solutions tend to negatively impact the chip performance. One solution which does not degrade the performance is to use advanced cooling material and innovative cooling technologies as reviewed next. Cooling solutions are generally classified as *passive* vs. *active cooling* [1]. Passive cooling techniques have no moving parts and do not need any power source to operate, whereas active cooling methods exploit moving parts and/or require an external power source. Reference [1] points out that a common disadvantage of

various passive and active techniques is their low heat-pumping capability [1]. In particular, none of the traditional techniques has the ability to pump heat fluxes higher than 1,000W/cm². Note also that active methods suffer from reliability issues and some of them, which provide a relatively high heat pumping capability (e.g., the direct jet impingement method), cannot be incorporated inside the chip package because of their large size. A new active cooling method called *thermoelectric cooling* has recently caught attention especially for cooling high-end multi-core processor chips [1].

Thermoelectric coolers (TECs) are active devices that work based on the *Peltier effect*. This effect allows the device to absorb the heat from one side and release it to the other side when electrical current passes through it. The amount of cooling is linearly proportional to the amount of driving current. Notable features of TECs are the following. 1) Compact size—TECs can be built as thin as tens of micrometers and their area can be smaller than 1mm². These devices have the right size to exclusively cover typical hot spots on a chip. 2) Fast response time—Thin-film TECs have very fast response times in the order of a few milliseconds. 3) High reliability—These devices have no moving parts, and hence, can last longer than other active cooling solutions. Commercial TECs are expected to work for more than 11 years [2]. 4) High controllability—TECs can be controlled at the granularity of fractions of a degree of Celsius and can cool down a chip below the ambient temperature. 5) Very high heat pumping rate—It has been shown that thin-film TECs can pump high heat fluxes as large as ~1,300W/cm² [3].

The unique features of TECs make them a perfect candidate for cooling a chip. Unfortunately, Joule heating occurs as an adverse phenomenon during the cooling process by TECs, which causes the device to dissipate heat when current flows through it. Both of the heat rejected from the hot spot and the heat generated by the TEC (as a result of Joule heating) must thus be disposed to the ambient; otherwise, the accumulated heat on the hot side of TEC adversely affects the performance of the TEC.

Apart from TECs, standard convention cooling techniques may be used: 1) natural-convection and 2) forced-convection. The first method is useful when the total amount of heat to be disposed is small. In the second method, forced-convection cooling allows more heat to be pumped from the chip using a fan. This extra ability comes at the cost of increased cooling power consumption. In this case, the total cooling power of the chip will be equal to the power usage of TECs and the fan. Moreover, simultaneously controlling the fan and TECs such that the entire system meets its thermal and power constraints is a challenging task. If TECs are driven by a high current level and the fan rotation speed is set to be too low, the rejected heat is trapped between the TEC and the fan, and hence, the hybrid cooling approach will not be effective. On the other hand, if the driving current of TECs is set to be too low but the fan rotation speed is set to be high, there is not enough pumped heat for the fan to blow away. Moreover, setting the fan speed and the TEC driving current to high levels increases the cooling power consumption, which negatively affects the power efficiency.

The argument presented in the previous paragraph suggests that TEC driving current and the fan rotation speed are two interrelated variables that directly affect the system temperature and the system cooling power consumption. As a result, there should be an optimum operating point at which the fan and TECs can work whereby the system thermal constraints will be met and the total cooling power will be minimized. In this paper, we focus on this joint optimization problem.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

DAC '14, June 01 - 05 2014, San Francisco, CA, USA
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-2730-5/14/06...\$15.00.
<http://dx.doi.org/10.1145/2593069.2593186>

To optimize TECs and fans, an accurate study of these devices is necessary. This study requires a simple, yet accurate, model of a hybrid chip cooling assembly in order to streamline the problem formulation and ease the process of finding the solution. An important consideration in this optimization problem is the leakage power, which is exponentially dependent on the temperature. Investing more power in the cooling may pay off well as a result of a dramatic power saving in the chip leakage power consumption. Therefore, the leakage power is considered in the proposed model and in the formulation.

Key contributions of this work are the following:

- 1) A compact thermal model for the hybrid chip cooling assembly comprised of TECs and a fan.
- 2) An optimization framework, called *OFTEC* (*optimization of forced-convection and thermoelectric coolers*), for minimizing the cooling-related power consumption while satisfying a maximum die temperature limit.

We show that OFTEC can meet the thermal constraints in all of the tested benchmarks whereas a system without TECs fails to meet the constraints in five out of eight benchmarks. In the remaining three benchmarks, OFTEC performs more power efficiently compared to a system without TECs by consuming 5.4% less power on average while keeping the hottest spot 3.7°C cooler on average. For all of the eight benchmarks, the average runtime of OFTEC is 437ms while the slowest runtime is 693ms. Moreover, it is shown that a system which adopts TECs as the only cooling method cannot avoid the thermal runaway situation in these benchmarks.

The rest of this paper is organized as follows. Section 2 reviews the principles of thermoelectric cooling. Section 3 describes the related work. Section 4 presents thermal models that are used in this paper. Next, Section 5 explains the problem formulation and the proposed solution. Section 6 presents experimental results and Section 7 concludes the paper.

2. PRINCIPLES OF THERMOELECTRIC COOLING

In this section, the key principles of thermoelectric cooling are reviewed. The presented equations are well known in the field of thermodynamics. Interested readers may refer to reference [1] for detailed discussions.

Thermoelectric coolers are compact devices which are made of pairs of N- and P-type semiconductor pellets. When current flows through a P-type pellet (from the positive terminal to the negative terminal), heat flows in the same direction, i.e., heat is absorbed from the positive side, which is called *cold side*, and released to the negative side, which is called *hot side*. The heat flow direction in an N-type pellet is the reverse of that of the P-type pellet. Usually several N-P pairs are connected electrically in series and thermally in parallel to increase the amount of heat rejection. Figure 1 shows a 3×3 array of TECs (a total of 9 N-P pairs).

The heat absorbed per unit time from the cold side and the hot side are denoted by \dot{q}_c and \dot{q}_h , respectively. They can be calculated as

$$\dot{q}_c = N \left(\alpha T_c I_{TEC} - K_{TEC} \Delta T - \frac{1}{2} R_{TEC} I_{TEC}^2 \right) \quad \text{and} \quad (1)$$

$$\dot{q}_h = N \left(\alpha T_h I_{TEC} - K_{TEC} \Delta T + \frac{1}{2} R_{TEC} I_{TEC}^2 \right), \quad (2)$$

where N is the number of TECs connected electrically in series; α is the *Seebeck coefficient*; T_c and T_h are the temperature of the cold side and hot side (in Kelvin), respectively; K_{TEC} is the thermal conductance of the TEC; ΔT is the temperature difference between the hot side and the cold side (i.e., $T_h - T_c$); R_{TEC} is the electrical resistance of the TEC; and I_{TEC} is the current flowing through the TEC. The first term in these equations captures the *Peltier effect* which is the cooling phenomenon, the second term signifies the heat conductivity of the TEC material, and the third term is the Joule heating effect. Note that the second and the third terms have adverse effects in cooling and hence have a negative sign. Moreover, the 1/2 coefficient for the Joule heating is added because it is approximated that half of the heat is released in the cold side and the other half is released in the hot side.




Figure 1. A 3×3 array of TECs

Also note that the Joule heating quadratically depends on the TEC driving current whereas the Peltier effect linearly depends on it. In Equations (1) and (2), the *Thomson effect* is not considered because of its negligible effect.

Power consumption of N TECs is equal to the difference between \dot{q}_h and \dot{q}_c and may be written as follows

$$P_{TEC} = \dot{q}_h - \dot{q}_c = N(\alpha \Delta T I_{TEC} + R_{TEC} I_{TEC}^2). \quad (3)$$

3. PREVIOUS WORK

Many studies have been conducted in the area of thermoelectric cooling. Most of them focus on improving the material that the device is made of and the manufacturing process. Reference [1] presents a comprehensive survey on TEC principles and the manufacturing advances in the recent years.

Recently, TEC has caught attention for cooling processor chips. Reference [4] uses TECs in order to cool down microprocessors in a datacenter and reduce the total cooling cost while maintaining high reliability. This paper mainly focuses on the steady-state analysis of TECs and uses a constant *coefficient of performance* (COP) for modeling TECs, which is defined as the ratio of the heat removed per second to the TEC power consumption. This method is inaccurate and too coarse grain. Moreover, the fan speed and its power consumption are assumed to be constant. Reference [5] shows the significance of the transient behavior of TECs in VLSI die cooling. It presents two simple controllers: *threshold based controller*, which turns on or off TECs when the temperature goes above or below a certain temperature, and *maximum cooling based controller*, which uses the hysteresis effect to decrease the number of ON/OFF transitions of TECs. In both controllers, TECs are supplied with a constant current to effect a state change. References [6] and [7] formulate the selective deployment of TECs on top of a chip in order to achieve the maximum cooling (lowest temperature). The motivation is that excessive deployment of TECs adversely affects the temperature of the device because of lateral heating among TECs. Moreover, deploying unnecessary TECs increases the power consumption of the cooling solution. The focus of these two papers is on the steady-state analysis of TECs. Reference [8] incorporates the leakage power into the TEC thermal model and suggests a new COP formulation for the entire cooling package. Then the optimum TEC current for achieving the minimum temperature or the maximum COP for the entire system is found through simulations for the entire chip package. The aforementioned papers in this paragraph do not account for any additional cooling method. As it will be shown in the experimental results section, TECs cannot sufficiently cool a very hot chip and avoid the thermal runaway situation. Hence, another cooling method or thermal management technique should be considered.

References [9] and [10] formulate the dynamic thermal management problem as a convex optimization in which the objective function is the total throughput of the system (which has to be maximized). Chip power consumption and die temperature are constraints of the problem formulation. Optimization variables are the frequency of CPU cores. Note that no active cooling technology is considered in these two papers. Reference [11] considers the fan speed, CPU frequency, and supply voltage as optimization variables in order to minimize the total energy consumption of the system. However, the thermoelectric cooling technique is not considered. Moreover, this paper assumes a lumped thermal model for a processor which sacrifices the accuracy of the model at the cost of a simplified model. Furthermore, this simplification may leave the hot spots on the chip since the lumped model considers the average temperature for the entire processor die.

4. MODELING

Figure 2 shows a typical cooling package assembly of a microprocessor in which TEC modules are incorporated. As can be seen, TECs are located between the *thermal interface material* (TIM) and the heat spreader for better heat conductivity. The heat spreader is also connected to the heat sink through another layer of TIM.




Figure 2. A chip assembly with its cooling solution

Using the duality between thermal and electrical phenomena, an electrical circuit model of heat flow in any physical system can be built [12]. An electrical circuit is convenient because it can be easily analyzed by using well-known circuit laws (such as KVL and KCL) and simulated with the aid of circuit simulators such as SPICE. To make this circuit model, each physical component is decomposed into several elements. Increasing the number of these elements increases the accuracy of the model; however, it also increases the complexity of the electrical circuit model, and thus, makes the analysis slow.

The processor package is comprised of seven layers: 1) PCB, 2) chip, 3) TIM1, 4) TEC, 5) heat spreader, 6) TIM2, 7) heat sink, and 8) fan. Layers 1, 3, 5, and 6 are referred to as $L_{conduct}$ in this paper. The elements in $L_{conduct}$ only conduct heat (i.e., they do not generate or absorb heat). Therefore, in the electrical circuit model, these elements are modeled as resistances as shown below.




Figure 3. An element in $L_{conduct}$ modeled by six resistors

Layer 2 is referred as L_{chip} which not only conducts heat similar to $L_{conduct}$, but also generates heat. The power consumption in this layer has two parts: *dynamic* and *leakage power*. Dynamic power is independent of the temperature and is not affected by the cooling solution. On the other hand, the leakage power depends exponentially on the temperature. In order to calculate the leakage power quickly, one may iteratively calculate it based on an initial temperature, update the temperature based on the calculated leakage power, and recalculate the leakage power again with the new temperature until the process converges. Reference [13] suggests to use the linear term of the Taylor series in the expansion of the leakage power equation. It is shown that this estimation speeds up the convergence dramatically. This linear estimation for element i can be written as follows:

$$p_{i,leakage} = a(T_i - T_{ref}) + b, \quad (4)$$

where a and b are the Taylor expansion coefficients, T_i is the temperature of element i , and T_{ref} is the reference temperature around which the Taylor series is expanded. This temperature is usually set as the average temperature of the chip or a particular functional unit inside the chip in order to increase the accuracy of the estimation, and consequently, speed up the aforesaid iterative method.

The next layer, the TEC layer, shows three different behaviors, namely heat absorption, heat rejection, and heat generation. Hence it is further broken into three sub-layers: $L_{TEC,Abs}$, $L_{TEC,Rej}$, and $L_{TEC,Gen}$. The power absorption, rejection, and generation for each element i in the abovementioned sub-layers can be calculated as follows:

$$p_i = -\alpha I_{TEC} T_i, \quad i \in L_{TEC,Abs} \quad (5)$$

$$p_i = \alpha I_{TEC} T_i, \quad i \in L_{TEC,Rej} \quad (6)$$

$$p_i = R_{TEC} I_{TEC}^2 + \alpha \Delta T I_{TEC}, \quad i \in L_{TEC,Gen} \quad (7)$$

where ΔT is the temperature difference between the upper (hot side) and the lower (cold side) elements. Note that Equation (7) is similar to

Equation (3) with the difference that it is written for one TEC. This equation defines the power consumption of a TEC unit. Other equations listed above, i.e., the power absorption and rejection, are for modeling and do not contribute to the power consumption of a TEC. Figure 4 depicts an electrical equivalent of a TEC for the steady-state analysis in which these sub-layers have been identified.

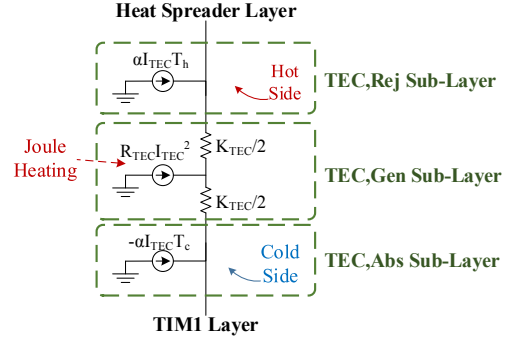


Figure 4. An electrical model for a TEC

Finally, layers 7 and 8 are categorized as $L_{HS\&fan}$. For the laminar airflow, the power consumption of a fan as a function of its rotation speed, ω , may be estimated as

$$P_{fan} = c\omega^3, \quad (8)$$

where c is a constant which depends on the air viscous friction, air density, and the radius of fan blades [14]. The thermal conductance of the heat sink depends on the air flow. A fan can change the air flow. Therefore, the collective thermal conductance of the heat sink and the fan together can be written as a function of ω . Using the calculation methodology employed in HotSpot 5 [12] and performing the curve fitting, the thermal conductance of the heat sink and the fan can be written as

$$g_{HS\&fan} = p \cdot \ln(q \cdot \omega) + r, \quad \omega \gg 1 \text{ rad/s}, \quad (9)$$

where p and r are fitting parameters, which depend on the material and physical properties of the heat sink, the fan, and air (such as air density and air thermal conductivity). Parameter q is added to make the logarithm value dimensionless so that both sides have the same unit dimension. In this paper, this value is simply considered as 1s and p and r are adjusted accordingly. For small values of ω , $g_{HS\&fan}$ can be estimated as the thermal conductance of heat sink (g_{HS}) which is constant under steady conditions of ambient air.

5. PROBLEM FORMULATION

5.1. Problem Statement

In this paper, the aim is to minimize the cooling power consumption of the entire chip package subject to the system thermal constraints. Since the leakage power is a function of temperature which is affected by the cooling efficiency, it is also included in the objective function. This problem can thus be formulated as shown in Optimization 1. (The optimization variables in this formulation are ω and I_{TEC} .)

$$\min_{\omega, I_{TEC}} \mathcal{P} := P_{leakage} + P_{TEC} + P_{fan} \quad (10)$$

where:

$$P_{leakage} = \sum_{i \in L_{chip}} p_{i,leakage} \quad (11)$$

$$P_{TEC} = \sum_{i \in L_{TEC,Gen}} p_i \quad (12)$$

$$P_{fan} = c\omega^3 \quad (13)$$

subject to:

$$\mathbf{G}(\omega) \vec{T} = \vec{P}(\omega, I_{TEC}) \quad (14)$$

$$T_i < T_{max}, \quad i \in L_{chip} \quad (15)$$

$$0 \leq \omega \leq \omega_{max} \quad (16)$$

$$0 \leq I_{TEC} \leq I_{TEC,max} \quad (17)$$

Optimization 1. Cooling power minimization subject to the thermal and the physical constraints

Equations (11)-(13) define the terms in the objective function. Equation (11) defines the leakage power as the sum of the leakage power of the elements ($p_{i,leakage}$) in the chip layer. Equation (12) expresses the TEC modules' power consumption as the sum over the power consumption of all TECs in the TEC generation sub-layer,

which was given in Equation (7). Equation (13), similar to Equation (8), defines the fan power consumption.

Next, the constraints are presented. Constraint (14) is a system of equations derived from KCL equations for all of the nodes in the equivalent electrical circuit; the total current (dual of power in the thermal model) that leaves a node (left hand side) is equal to the current (power) that enters a node (right hand side). Matrix \mathbf{G} is defined as follows:

$$\mathbf{G} = \begin{bmatrix} \sum g_{1,i} & -g_{1,2} & \dots & -g_{1,n} \\ -g_{2,1} & \sum g_{2,i} & \dots & -g_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ -g_{n,1} & -g_{n,2} & \dots & \sum g_{n,i} \end{bmatrix} \quad (18)$$

where $g_{i,j} = g_{j,i}$ is the thermal conductance between element i and element j . All of these values are constant except the thermal conductance between the heat sink/fan and the ambient, which is equal to $g_{HS\&fan}$. This value is a function of ω as described in the previous section, and hence, matrix \mathbf{G} is a function of ω . Vector \vec{T} keeps the temperature of all elements in the thermal model where the temperature of element i is denoted by T_i . Vector \vec{P} contains the power consumption of all elements in all of the layers in the thermal model, where the power consumption of element i is denoted by p_i . The definition of p_i values for elements of each layer is presented in the previous section. Note that the dynamic power consumption of the elements in the chip layer is considered as an input to the problem.

Vector \vec{P} is a function of both of the optimization variables, i.e., I_{TEC} and ω . Here, it can be seen that using a linear estimation for the leakage power as opposed to a constant value does not add any computational complexity to constraint (14), since it is already a system of linear equations with respect to T_i values. As it is explained earlier, this estimation speeds up the leakage power calculation.

Constraint (15) ensures that the temperature of all elements in the chip layer remains below a certain threshold (T_{max}). Constraint (16) and (17) enforce physical constraints. More precisely, constraint (16) sets an upper bound (ω_{max}) and a lower bound (zero) for the rotational speed of the fan, whereas constraint (17) imposes a lower bound (zero) and an upper bound on the driving current of TECs ($I_{TEC,max}$). Note that if the TEC current exceeds a certain threshold, the TEC will be damaged.

5.2. Proposed Solution

The problem formulation presented in Optimization 1 is not convex. Moreover, due to the iterations required for the calculation of leakage power, the objective function \mathcal{P} can only be determined numerically for a given ω and I_{TEC} . This problem is classified as a *constrained nonlinear program* (CNLP). We experimented with three state-of-the-art nonlinear optimization techniques for solving this problem, namely, *interior-point* method, *trust-region* technique, and *active-set sequential quadratic programming* (SQP) method [15]. It turns out that the last technique, i.e., the active-set SQP method performs the best for our formulation both in terms of solution quality and speed. This technique is briefly explained next.

The active-set SQP tries to find a solution for the *Karush-Kuhn-Tucker* (KKT) conditions, which are necessary conditions for the optimality of a solution. At an optimum point, when a constraint is active, its contour is tangential to that of the objective function. This means that the gradient of the objective function is equal to the gradient of the active constraint, though it may have different absolute value. *Lagrangian multipliers* are used to compensate for different gradient sizes. These multipliers are non-zero when a non-equality constraint is active and zero otherwise. The active-set SQP method tries to solve the KKT conditions iteratively by approximating those using convex quadratic programs (QP). Solving QPs allows determining the search direction. Having the search direction and the step length (which can be found through a *line search* method), a near-optimal solution can be found. Since the non-convexity of the optimization function of our interest is minor (this will be shown in the experimental results), the active-set SQP method produces high quality results very quickly. A

detailed explanation of the active-set SQP method can be found in [15].

The active-set SQP method, similar to other nonlinear optimization techniques, requires an initial feasible solution. Finding an initial solution is not trivial since the relation between constraint (15) and optimization variables is set by the set of nonlinear equations listed in constraint (14). On the other hand, minimization of the objective function irrespective of the constraints may violate constraint (15). To address this difficulty, a new optimization problem is formulated in order to find an initial feasible solution for the original problem. The formulation is listed below. (In this optimization problem, similar to the previous one, the optimization variables are ω and I_{TEC} .)

$$\min_{\omega, I_{TEC}} \mathcal{J} := \max_{i \in L_{chip}} \{T_i\} \quad (19)$$

subject to:

$$\mathbf{G}(\omega)\vec{T} = \vec{P}(\omega, I_{TEC}) \quad (20)$$

$$0 \leq \omega \leq \omega_{max} \quad (21)$$

$$0 \leq I_{TEC} \leq I_{TEC,max} \quad (22)$$

Optimization 2. Minimizing the maximum chip temperature subject to the thermal and the physical constraints

Finding an initial solution for this problem is trivial; it can be done by arbitrarily selecting (ω, I_{TEC}) such that the pair satisfies domain constraints (21) and (22). In this paper, we simply set these values to $(\frac{\omega_{max}}{2}, \frac{I_{TEC,max}}{2})$. Constraint (20) will not be violated since T_i 's will be adjusted according to the selected (ω, I_{TEC}) . Optimization 2 is an interesting problem by itself since it minimizes the maximum die temperature, which leads to the minimization of the maximum leakage power consumption of elements and also it slows down the aging rate of the element on the chip layer with the highest temperature. So this solution has its own applications as long as the cooling power consumption is not a concern. If it turns out that the minimized $\max_{i \in L_{chip}} \{T_i\}$ is greater than T_{max} , it can be concluded that Optimization 1 has no solution, i.e., it is *infeasible*. Moreover, the solver can stop the optimization procedure as soon as it finds the first solution which makes $\max_{i \in L_{chip}} \{T_i\}$ smaller than T_{max} . Having an initial feasible solution for Optimization 1, one can use the active-set SQP method to approximately solve it. Algorithm 1, called *optimization of forced-convection and thermoelectric coolers* (OFTEC), presents this approach.

Algorithm 1: OFTEC

Input: Physical characteristics of the cooling package and the dynamic power consumption of each element in the chip layer.

Output: ω^* and I_{TEC}^*

1. $(\omega_0, I_{TEC,0}) \leftarrow (\frac{\omega_{max}}{2}, \frac{I_{TEC,max}}{2})$

2. **If** $\mathcal{J}(\omega_1, I_{TEC,1}) > T_{max}$ **then**

3. $(\omega_1, I_{TEC,1}) \leftarrow$ **Call** the active-set SQP method to solve Optimization 2 with the initial solution $(\omega_0, I_{TEC,0})$. Stop the optimization whenever $\mathcal{J}(\omega, I_{TEC}) < T_{max}$.

4. **If** $\mathcal{J}(\omega_1, I_{TEC,1}) > T_{max}$ **then**

5. **Return failed** //no solution is found

6. $(\omega^*, I_{TEC}^*) \leftarrow$ **Call** the active-set SQP method to solve Optimization 1 with the initial solution $(\omega_1, I_{TEC,1})$.

7. **Return** (ω^*, I_{TEC}^*)

6. EXPERIMENTAL RESULTS

6.1. Simulation Setup

In order to evaluate OFTEC, the flow shown in Figure 5 is developed. The experiments target the Alpha 21264 processor. PTscalar [16] is used as the performance/power simulator in order to generate the dynamic power trace for the benchmarks which are selected from the MiBench benchmark suite [17]. The maximum power consumption for each element in the chip layer is selected to be passed to OFTEC along with the cooling package configuration and the chip floorplan so that it finds the near-optimum I_{TEC} and ω . Note that this flow is not limited to the aforementioned selections of the processor and

performance/power simulators; any other set of processor and simulators can be used.




Figure 5. The evaluation flow for OFTEC

The active-set SQP method is implemented in MATLAB to solve the non-convex optimizations. The value of two objective functions presented in the previous section (i.e., \mathcal{T} and \mathcal{P}) are calculated numerically through a thermal simulator given I_{TEC} and ω . This simulator is a modification of *Teculator* [8] in which the models presented in Section 4 are incorporated. Note that this simulator performs no optimization. In order to streamline the connection between the simulator (which is written in the C language) and the MATLAB code, *Teculator* is compiled with MATLAB MEX compiler. This gives two important advantages. First, the code can be reused with a minor change, i.e., only an interface between the simulator and the MATLAB code should be implemented. Second it does not degrade the performance of the C code whereas re-implementing the simulator in MATLAB would dramatically affect it.

Based on the experiments presented in reference [11], the fan power constant c in Equation (8) is estimated as $1.6 \times 10^{-7} \text{ J}\cdot\text{s}^2$. Moreover, ω_{max} , $I_{TEC,max}$, and T_{max} are set to 524 rad/s (=5000RPM), 5A, and 90 °C (363K), respectively. The ambient temperature of the chip is assumed as 45 °C (318K). The processor package assembly used for simulations has a similar configuration to Figure 2. Table 1 shows dimensions and thermal conductivity of all layers used in the simulations (except the TEC sub-layers which are taken from reference [8]). The entire surface of the processor is tiled with TECs except the instruction and data caches which are remained uncovered since they do not show any hot spots in the experiments. This observation agrees with the results presented in reference [12]. Moreover, avoiding the excessive deployment of TECs helps eliminate the power they are consuming and heating their neighbor TECs [6][7]. The deployed TECs are connected electrically in series and driven by the same current value.

Table 1. Thermal conductivity and dimensions of various layers in the chip package

Layer	Thermal Conductivity (W/(m·K))	Dimensions
Chip	100	15.9mm×15.9mm×15μm
TIM 1	1.75	15.9mm×15.9mm×20μm
Heat spreader	400	30mm×30mm×1mm
TIM 2	1.75	30mm×30mm×20μm
Heat sink	400	60mm×60mm×7mm

McPat [18] is used in order to estimate the leakage power of the Alpha 21264 processor (whose model comes with the tool) using the 22nm CMOS process technology. The simulation is done for ten temperature values distributed evenly in the range of 300K to 390K. Using these ten values as T_{ref} in Equation (4), Taylor expansion coefficients a and b are calculated by performing linear regression. Moreover, p and r in Equation (9) are set to 0.97W/(m·K) and -0.25W/(m·K), respectively. g_{HS} is also set to 0.525W/(m·K).

We considered two systems as baseline for our comparisons: 1) A system without any TECs equipped with a fan controlled by variable speed. The speed is set using a method similar to OFTEC with the difference that no TEC current is required to be found. 2) A system with a fan with fixed rotation speed where $\omega = 2000\text{RPM}$. In our experiments, unlike OFTEC which utilizes TECs, we realized that the two baselines fail in all except one of the benchmarks to avoid the thermal runaway phenomenon. The reason is that the thermal conductivity of the material that TECs are built from is much higher than that of common thermal pastes used in the TIM1 layer [3]. When TECs are deployed, they are placed on top of the TIM 1 layer (see Figure 2), which results in increasing the overall thermal conductivity

of the cooling package compared to the case without TECs. However, the passive use of TECs is not common because thermal pastes with high heat conductivity are cheaper than TECs. So to make the comparison fair, the conductivity of the TIM1 layer in the baselines is set equal to the overall conductivity of TIM1 plus the TEC.

6.2. Simulation Results

Figure 6 (a) and (b) show the objective function of two optimization problems drawn for different values of ω and I_{TEC} . These figures belong to the *Basicmath* benchmark. Objective functions of other benchmarks generally have the same shape but they are not presented here due to the lack of space. As can be seen, both functions have a smooth and convex shape; however, some minor non-convexities exist. Since the size of these non-convexities is small, the active-set SQP can find a very high quality solution.

It is important to note that the value of \mathcal{P} and \mathcal{T} tends to infinity for small values of ω . This is shown in the figures by dark red color. The physical interpretation is that due to the lack of enough cooling, the system traps in a *thermal runaway* situation where the high leakage power causes the temperature to increase, and the elevated temperature increases the leakage power further. This cycle eventually ends in a burned chip. As can be seen, increasing I_{TEC} alone cannot rescue the chip from the thermal runaway situation; ω should also be increased to about 150 RPM at the same time. This signifies the motivation of the paper that TECs cannot pump the heat effectively without further assistance of other cooling techniques to dispose the extracted heat. Also note that the minimum of the two objective functions occur at different points which shows the importance of each of the optimization problems. In fact, the surface chart shown in Figure 6(a) is the thermal constraint of the Optimization 1 and its objective function is depicted in Figure 6 (b). In Figure 6 (a), the minimum occurs at almost the middle of the $(\omega-I_{TEC})$ plane. That is why in the first line of Algorithm 1, we set the initial value of (ω, I_{TEC}) as $(\frac{\omega_{max}}{2}, \frac{I_{TEC,max}}{2})$. Further increase of ω and I_{TEC} values causes more heat to be generated by the fan and TECs than the cooling they provide. On the other hand, in Figure 6 (b), the minimum occurs near the origin.

Figure 6 (c) and (d) show the results of performing Optimization 2 (i.e., line 3 in Algorithm 1). Figure 6 (c) depicts the maximum die temperature (\mathcal{T}) achieved by OFTEC and two baselines. The thermal threshold (T_{max}) is shown by a dashed line in this figure. As can be seen, OFTEC could meet the thermal constraint in all benchmarks, whereas two baselines failed to cool down the system in five benchmarks which are identified by a red dashed box. These five cases should be further cooled down using other thermal management techniques such as reducing the voltage/frequency of the chip or throttling different functional units which leads to performance degradation. Moreover, on average, OFTEC could achieve more than 13°C lower temperature compared to the other two cases. Figure 6 (d) compares the power consumption of these three methods. As can be seen, OFTEC has the highest power consumption when the objective function is to minimize the maximum temperature. This extra power is consumed mostly by TECs.

Figure 6 (e) and (f) show the results of performing Optimization 1. Results of two baselines are omitted in five benchmarks since they could not meet thermal constraints and hence do not provide meaningful results. In Optimization 1, OFTEC addresses the trade-off between the cooling power consumption and the maximum chip temperature. Figure 6 (e) shows that OFTEC slightly increases the temperature in order to reduce the cooling power consumption. Note this increase is done such that the system temperature still meets thermal constraints. Figure 6 (f) compares the power consumption of these three methods. OFTEC has the minimum power consumption among three cooling methods. In the comparable cases, in which all of them could meet the threshold, OFTEC could save 0.35W and 1.04W (or 2.6% and 8.1%) on average compared to the variable and fixed ω methods, respectively. OFTEC could achieve these results by keeping the highest chip temperature cooler by 3.7°C and 3.0°C than the variable and fixed ω methods, respectively. This chart clearly shows that OFTEC only allow necessary cooling power to be dissipated in

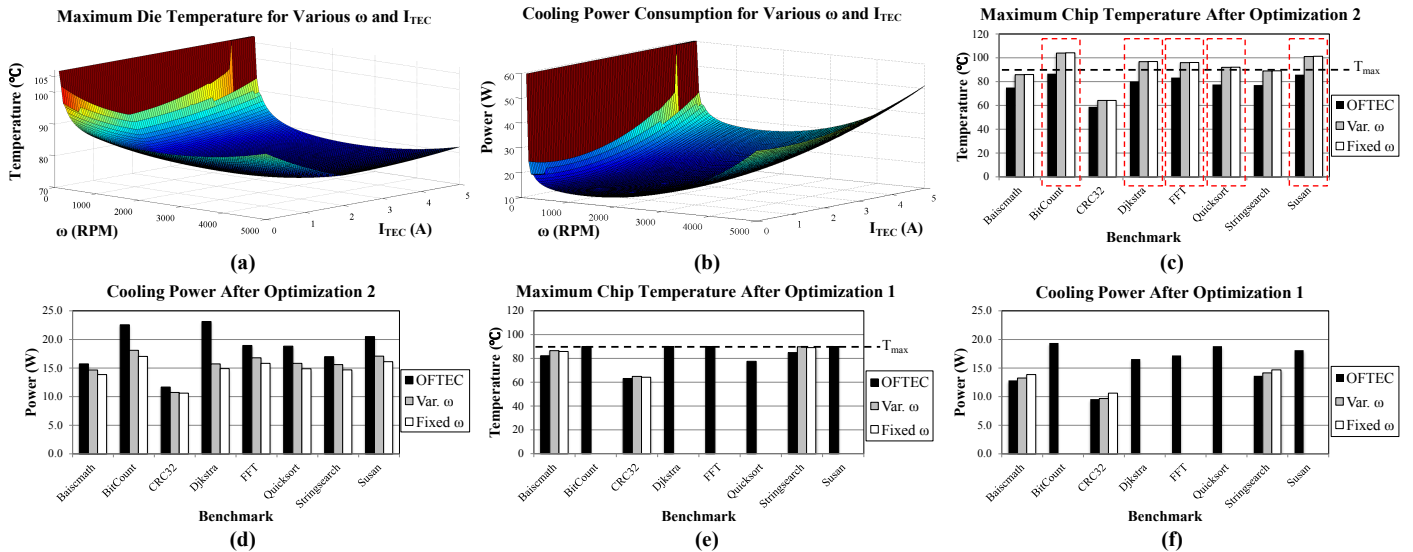


Figure 6. (a) Objective function of Optimization 1 and (b) Optimization 2 for the Basicmath benchmark. (c) Comparison among the maximum chip temperature achieved by three methods in Optimization 2 and (e) Optimization 1. (d) Comparison among the cooling power consumption achieved by three methods in Optimization 2 and (f) Optimization 1.

order to meet thermal thresholds. If thermal constraints can be met with lower power, OFTEC adjusts I_{TEC} and ω accordingly.

Table 2 shows results that OFTEC could produce for eight MiBench benchmarks and their respective runtimes on a system with an Intel Core i7-3770 CPU (running at 3.4GHz) and 8GB memory. As can be seen, I_{TEC}^* and ω^* values are increased when the input dynamic power is high and more cooling is required to cool down the chip. Moreover, OFTEC is a fast algorithm which can find the solution in 437ms on average.

Table 2. Results of OFTEC for MiBench benchmarks

Benchmark	I_{TEC}^* (A)	ω^* (RPM)	Runtime (ms)
Baiscmath	0.68	1352	426
BitCount	2.30	2451	693
CRC32	0.37	1114	239
Dijkstra	1.14	2516	430
FFT	0.99	2490	353
Quicksort	2.83	2433	385
Stringsearch	0.74	1399	278
Susan	1.81	2509	690

We expect that implementing the active-set SQP method in C language will substantially speed-up the runtime which allows OFTEC to be used as an online controlling algorithm. Also, with the current runtime of OFTEC, one can classify the input dynamic power vector to different categories and pre-calculate optimization solutions and store them in a look-up table. In this way, the desired controlling values can be accessed immediately. Moreover, TECs can improve the heat removal capacity of steady state cooling solutions for a short period of time (i.e., order of a second). This is because the Joule heating effect appears with a delay but the Peltier effect shows up immediately [8]. This phenomenon can be used before results of OFTEC become ready. Reference [8] suggests to increase I_{TEC} (i.e., the optimum TEC current) by about 1A for 1s to reap the benefit of transient cooling. These approaches will be investigated as future extension of this work.

7. CONCLUSION

This paper presented a thermal model for a hybrid cooling assembly comprised of TECs and a fan. Then a formulation for the minimum cooling power optimization problem subject to the system thermal and physical constraints was proposed in which optimization variables were the TEC driving current and the fan speed. Next, an optimization framework called OFTEC was developed in order to solve this problem. Simulation results showed that OFTEC can meet thermal constraints in all of benchmarks whereas a system without TECs fails to meet the constraints in five out of eight benchmarks. In the remaining three benchmarks, OFTEC performed more power efficiently compared to a system without TECs by consuming 5.4%

less power on average while keeping the hottest spot 3.7°C cooler on average. For all of the eight benchmarks, the average runtime of OFTEC was 437ms. Moreover, it was shown that a system which adopts TECs as the only cooling method cannot avoid the thermal runaway situation in these benchmarks.

8. ACKNOWLEDGEMENTS

This research was sponsored in part by grants from the NSF SHF and the DARPA PERFECT programs.

9. REFERENCES

- A. Bar-Cohen and P. Wang, "On-Chip Thermal Management and Hot-Spot Remediation," in *Nano-Bio- Electronic, Photonic and MEMS Packaging*, Springer, 2010.
- "Tellurex - An introduction to thermoelectrics." [Online]. Available: <http://www.tellurex.com/technology/design-manual.php>.
- I. Chowdhury et al., "On-chip cooling by superlattice-based thin-film thermoelectrics," *Nat. Nanotechnol.*, vol. 4, no. 4, pp. 235–238, 2009.
- S. Biswas et al., "Fighting fire with fire: modeling the datacenter-scale effects of targeted superlattice thermal management," in *ISCA*, 2011.
- B. Alexandrov et al., "Prospects of active cooling with integrated superlattice based thin-film thermoelectric devices for mitigating hotspot challenges in microprocessors," in *ASP-DAC*, 2012.
- J. Long and S. O. Memik, "A framework for optimizing thermoelectric active cooling systems," in *DAC*, 2010.
- J. Long et al., "Optimization of an on-chip active cooling system based on thin-film thermoelectric coolers," in *DATE*, 2010.
- M. J. Dousti and M. Pedram, "Platform-Dependent, Leakage-Aware Control of the Driving Current of Embedded Thermoelectric Coolers," in *ISLPEd*, 2013.
- S. Murali et al., "Temperature-aware processor frequency assignment for MPSoCs using convex optimization," in *CODES+ISSS*, 2007.
- "Temperature Control of High-Performance Multi-core Platforms Using Convex Optimization," in *DATE*, 2008.
- D. Shin et al., "Energy-Optimal Dynamic Thermal Management: Computation and Cooling Power Co-Optimization," *IEEE Trans. Ind. Inform.*, vol. 6, no. 3, pp. 340–351, 2010.
- K. Skadron et al., "Temperature-aware microarchitecture," in *ISCA*, 2003.
- Y. Liu et al., "Accurate Temperature-Dependent Integrated Circuit Leakage Power Estimation is Easy," in *DATE*, 2007.
- I. Sato et al., "Characteristics of heat transfer in small disk enclosures at high rotation speeds," *IEEE Trans. Compon. Hybrids Manuf. Technol.*, vol. 13, no. 4, pp. 1006–1011, 1990.
- J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 2006.
- "PTscalar." [Online]. Available: <http://eda.ee.ucla.edu/PTscalar/>.
- "MiBench." [Online]. Available: <http://www.eecs.umich.edu/mibench/>.
- S. Li et al., "McPAT: an integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Micro*, 2009.