DESIGN METHODOLOGY FOR THERMAL MANAGEMENT USING EMBEDDED THERMOELECTRIC DEVICES

A Dissertation
Presented to
The Academic Faculty

By

Borislav P. Alexandrov

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy
in
Electrical and Computer Engineering

School of Electrical and Computer Engineering
Georgia Institute of Technology
December 2015

DESIGN METHODOLOGY FOR THERMAL MANAGEMENT USING EMBEDDED THERMOELECTRIC DEVICES

Approved by:

Dr. Saibal Mukhopadhyay, Advisor
Associate Professor,
School of ECE
Georgia Institute of Technology

Dr. Arijit Raychowdhury
Associate Professor,
School of ECE
Georgia Institute of Technology

Dr. Sudhakar Yalamanchili
Joseph M. Pettit Chair Professor,
School of ECE
Georgia Institute of Technology

Dr. Sung Kyu Lim
Dan Fielder Endowed Chair Professor,
School of ECE
Georgia Institute of Technology

Dr. Satish Kumar
Assistant Professor,
School of Mechanical Engineering
Georgia Institute of Technology

Date Approved: August 18, 2015

ACKNOWLEDGMENT

I would like to first express my deepest thanks to my parents, Siyka and Peter, for their love, support, and extreme sacrifice. Immigrating to a new country is not an easy task and my parents were able to leave their life behind in hopes of a better future for my brother and I. They showed me how to persevere and push on even in the darkest and most difficult of times. Without them I would not be the man I am today and I want to thank them for providing me with a chance to succeed. I also want to thank them for encouraging me to push on during my PhD even in times when I was down, but I mostly want to thank them for leading by example. Despite coming to a new country and adapting to a new culture, they never ceased their intellectual curiosity and continued with their studies. They showed me that success can be attained as long as you are willing to put your mind to it and work hard. I also want to thank my brother George for all the childhood memories, continuous support, and time spent together. He is the best younger brother one could ask for. I also want to thank all of my extended family in Bulgaria, who have always supported me and have spent countless hours on the phone and on Skype to make sure that I was always doing well.

I would like to acknowledge and express the highest gratitude to my advisor Dr. Saibal Mukhopadhyay for his support, guidance, and most importantly for letting me part of his lab throughout my PhD. From the first day of starting this long journey I knew that Dr. Mukhopadhyay was going to be a great advisor. Taking his course in my first semester was a wonderful experience and a chance for me to get to know him. I knew instantly that he would be a great advisor and that I would be able to learn lots from him. When I joined the GREEN Lab these thoughts became a reality. Dr. Mukhopadhyay was able to get me out of my comfort zone and inspire me to push myself further than I thought I could. His patience and encouragement towards simulations, testing chips, and paper submissions were tremendous. Even in times when I was down and frustrated, he was able to encourage

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

Modern integrated circuits and microprocessors continue to pack more transistors, increasing the power densities and increasing the chip temperatures. Higher chip temperature can severely limit system performance and force cores to operate at suboptimal power due to the inability to remove the heat and operate under the chip thermal limit. Modern systems take advantage of the thermal time constants of the packages to exceed the thermal design power (TDP) for a short duration. This work investigates the feasibility of using thin-film quantum dot thermoelectric materials to provide on-demand active cooling and thermal management. The aim is reduce the hot spot and overall chip temperature and improve performance by allowing the system to operate at higher power modes. Thermoelectric coolers are lightweight, have no moving parts, have high reliability, and can be integrated directly onto a heat spreader, making them attractive candidates for both high-performance and mobile devices. The research presented in this thesis studies the feasibility of using thermoelectrics for thermal management and develops a design methodology that can successfully reduce the thermal events experienced in integrated circuits by using embedded thermoelectric devices. This thesis develops a compact thermal model and simulation framework. The framework is used to study the system level prospects of any system and package with TEC assisted on-demand cooling considering thermal benefits, performance implications, and energy overheads. The thesis also develops control methods to efficiently control the TEC and implements the controllers in modern IC processes to assess their performance and overheads. The controllers are simulated within an expanded electro-thermal co-simulation environment as well as verified through experimental design in an IBM 130nm CMOS process. Finally the thesis presents a method to increase the energy efficiency of TEC assisted cooling by harvesting energy from the wasted heat from the chip during thermally non-critical events. The proposed approach is experimentally characterized using a test-chip designed in 130nm CMOS process with an external TE device.

CHAPTER 1

INTRODUCTION

For the last 50 years, the semiconductor industry has been driven by Moore's Law. This simple economic law observed and predicted by Gordon Moore and following the guidelines first suggested by Bob Dennard in 1974 [2], has driven the semiconductor industry to great lengths by doubling the number of transistors on a chip by a factor of 2 every 18-24 months. The reduction in chip area and increase in performance of transistors have led to powerful microprocessors and smartphones that have completely revolutionized society as it is today. More importantly because many more transistors could be packed into the same area, the cost of each transistor has exponentially declined. The scaling however has not been easy, and the industry has heavily relied on material innovation and the changing of the transistor structure. The first major challenge faced was the scaling of the gate insulator thickness, as once the gates got thinner, the probability of electrons tunneling directly across began to increase. High-k dielectric materials were therefore introduced that provided the transistor to maintain the same electrostatic control with a much thicker insulator that prevented the electrons from tunneling across. Another defining innovation that has further extended transistor scaling was the introduction of the 3-D transistor or FinFET. This extends the channel of the transistor into the 3$^{\text{rd}}$ dimension and allows for better controllability. This has also allowed for lower supply voltages of operation and the reduction of short channel effects. Today's 14nm transistors that will go into the latest chips are truly a remarkable piece of engineering.

Designers have also taken advantage of the large number of transistors as well as the faster transistors introduced in every generation. Packing more devices into a chip and increasing the frequency of operation has unfortunately led to a significant increase in power consumption. As the power of the chips began to exceed the capability of cost-effective

cooling solutions, designers had to reduce the frequency of operation as well as the oper-ating voltage of the chips. This led to performance reductions and a power wall due to the inability to scale voltage down further. Despite the restrictions, the increased performance requirements continued to push designers to design multi-core chips, effectively spreading the area over which the work was done and reducing the power density. Clever techniques like Intel's Turbo-Boosting were introduced that exceeded the allowable power of the chip for a short time in order to complete a task before the system reaches a thermal limit. As the need for increased performance is steadily continuing, new methods of thermal manage-ment are required. Recent advances in thin-film quantum dot thermoelectric materials have improved the figure of merit of thermoelectric coolers (TEC), and experimental integration within ICs have made them a strong candidate for thermal management. Thermoelectric coolers are lightweight, have no moving parts, have high reliability, and can be integrated directly onto a heat spreader, making them attractive candidates for both high-performance and mobile devices. The proposed research aims to study the feasibility of using thermo-electrics for thermal management and develop a design methodology that can successfully reduce the thermal events experienced in microprocessors and mobile devices by using embedded thermoelectric devices.

Chapter 2 provides a literature review of the existing thermal management technolo-gies and architectural methods proposed to mititgate thermal issues of intergrated circuits. The existing work in using TEC for thermal management are reviewed and the need for additional research is highlighted.

Chapter 3 develops the simulation framework for thermal characterization. We first develop a compact model of the TEC in SPICE and develop a 3D distributed RC grid to include the full chip and package to allow for rapid simulation and prototyping. The TEC model is fully validated to experimental measurements in the literature [3] as well and FVM computational models, providing up to 4X speedup in simulation time at less than 5% error. The TEC is integrated into a package and simulated to understand the thermal

performance characteristics. The package is simulated with the transient effects accounted for and specific heat capacity for each material. The transient simulations of the TEC and package show that TEC can significantly reduce transient fluctuations in chip temperature and due to the proximity to the silicon layer can provide on-demand cooling. The TEC is shown to provide up to 10°C of instant cooling and can be used to extend high power events and increase performance. Case studies are presented for both single core applications as well as multi-core scenarios. In the single core space we show to use a TEC to minimize the chip temperature during a known-length power event like turbo boosting, and discuss the energy implications of the TEC. In the multi-core study we show how to use the TEC to manage the temperature on chip for an infinitely long workload with core hopping. We show that using the TEC can allow the chip to sustain up to 20% additional power without reaching its thermal limit. The package and TEC model can be extended to any workload and system and can be used to provide direct feedback to micro-architecture development.

Chapter 4 uses the modeling framework and builds on it to develop a general co-simulation framework to simulate thermal and electrical characteristics simultaneously. This allows for rapid design of control methods to activate/deactivate the TEC and provide on-demand cooling. We develop the controllers in a 130nm IBM process and simulate them within the framework. The first control method is a threshold based controller (ThBC) that tries to maintain the chip temperature near a reference level for as long as possible. The second control method, MCBC, takes advantage of the initial cooling dip of a TEC upon turn-on and minimizes the average temperature during a transient high power event. Both control methods provide similar workload extension times of around 100ms. This allows a core to complete a workload without the need for dynamic thermal management techniques. Finally the control methods are integrated within a real processor core and verified for functionality with the latest deep sub-micron technology and benchmarks.

Finally the thesis presents a fully on-chip autonomous energy management system in chapter 5 to take advantage of the TEC by also using it in its reserve energy harvesting

mode. This system relies on the fact that thermal events occur with low probability within a system and the TEC is mostly idle. By taking advantage of idle power modes on chip we can use the TEC in reverse to harvest energy, as a TEG. This allows us to improve system efficiency by using the otherwise wasted heat energy. We develop a chip level implementation of the system which is fabricated in a 130nm IBM CMOS process from MOSIS. The system stores energy during thermally non-critical events and uses a boost regulator to harvest the energy into a capacitor. During high power events we use the stored energy to power the TEC current. Once depleted, we use the chip's voltage supplies to provide cooling energy. The test chip also verifies the TEC controls that were developed in chapter 4. The chip can be programmed to change the TEC current level and allow it to be used across multiple TEC materials.

CHAPTER 2

LITERATURE SURVEY

2.1   Power Dissipation Issues

The technology scaling of Moore's Law has enabled designers to pack more and more transistors in every new generation of a chip, pushing performance and enabling faster and cheaper computing. Unfortunately due to increases in the operating frequency from the faster transistor as well as larger active capacitances, packing more and more transistors into a small area increases the power density of the chip [4]. An early method to keep the power increases down was to reduce the supply voltage, however that traded off the maximum frequency (i.e performance) that the chip could operate at [5]. Despite this, supply voltages have continued to be scaled down further, reaching a limit imposed by the threshold voltage of the transistor [6]. In addition, with each generation, leakage currents [7] have also began to contribute significant power dissipation resulting in up to 50% of the chip power being dissipated as leakage. Despite the latest technology introduction of the 22nm FinFET [8], which reduces leakage by the introduction of a 3D gate with better short channel effect controllability, leakage power in chips is still large. Today's latest generation microprocessors [9] still have large leakage currents that increase the total chip power dissipation and cause significant idle mode power dissipation. Higher temperatures generated by packing higher power transistors into a smaller area can lead to performance limits imposed by the maximum temperature that the chip can sustain. In addition, localized high heat fluxes within a chip result in localized high temperature regions or hotspots [10, 11].

2.2   Architectural Techniques for Power Reduction

In order to deal with increasing power and temperatures generated across chips, various architectural methods have been developed. The most basic method, which is implemented in nearly every microprocessor today, is fan speed control. As the temperature of the chip

increases, the fan speed can be increased to provide higher convective cooling. The speed can be changed during processor operation and allows for simple run-time control of the temperature. However, some workloads can still generate significant power that causes the temperature to increase beyond fan speed controllability. This temperature limit, or thermal design power (TDP), can severely limit single threaded performance. Multi-core processors were later introduced [12] in order to increase the performance of systems by exploiting parallelism. An added benefit was the ability for improved thermal performance, as workload could be shifted between processors by thread migration. Chaparro et.al [13] have investigated thread migration in 16-core systems in order to balance the thermal performance of the chip by migrating high power workloads between cores. Others have explored optimization techniques for migration with thermal balancing policies and predictive workloads [14, 15]. Lastly Dynamic Voltage Frequency Scaling (DFVS) has been proposed [16, 17] to better control temperature fluctuations by balancing performance and the thermal response of the system.

2.3   Cooling Technologies

The architectural methods from the preceding section can only mitigate the thermal problem to a small degree, and in order to reduce the temperature of chips further, material advances in cooling technologies are required. Various technologies have been explored with their effects thoroughly studied. Dry phase change materials (PCM) have been explored [18], as a method of mitigating the temperature variations when embedded in a cooling solution within a chip. Essentially the material is embedded in the package solution and when the chip heats it up to the phase change point, the material will remain at the specific temperature until the duration of the phase change. Current materials have shown experimental integration with heat sinks [19] (PCM above the heat sink) for use in thermal management. Although the PCM is able to provide transient cooling at a certain temperature, the solution requires specific material engineering for the application so that the target

temperatures can be set. In addition, since the phase changes happen on time scales on the order of minutes, this cooling method is not very fast and cannot be controlled by the user. Liquid cooling has also been explored. In traditional heat-sink based solutions, the liquid is pumped around the heat sink taking away more heat than the air would be able to in itself. This is a rather expensive solution that requires tubes and a liquid storage device, but it is very effective at providing cooling to high performance systems. This is a limiting technology for mobile devices, where liquids and pumping systems are difficult to embed. More recently microfluidic liquid cooling solutions [20] have been integrated in 3D IC's that have shown tremendous heat removal capabilities even in the 3D stacking of 2 microprocessors. This a rather expensive solution however, as added steps are needed in a process in order to be able to drill the channels within the Si and metal layers, which can be prohibitively expensive. Recently, thermoelectrics have also been proposed for use in cooling of integrated circuits [3, 21], specifically for hot-spot mitigation. Thermoelectric coolers (TEC) are 2 terminal devices that operate on the principle of the Peltier effect by dumping heat from the hot side to the cold side when an electric current flows through the device. TECs are lightweight, have no moving parts, have high reliability, and can be integrated directly onto a heat spreader, making them attractive candidates for thermal management. The amount of heat taken away is proportional to the difference in the temperature and current through the TEC. The high figure of merit (ZT) of these thin film super-lattices and their high heat-flux pumping capability make them very attractive material for use in electronic devices. ZT is defined as a dimensionless unit that is used to compare thermoelectric materials. $ZT = S^2 \sigma T / \kappa$, where S is the Seebeck coefficient, T is the absolute temperature, and σ and κ are the electrical and thermal conductivities respectively. The higher the ZT the better the material is said to be when used as a TEC. Experimentally, integration of TECs with the thermal interface material (TIM) in a processor package has been demonstrated with total hotspot cooling of up to 15°C [21]. Further, as cooling is controlled by an electrical current, the TEC provides the opportunity for on-demand run-time control of this additional

cooling [3, 22].

2.4   TEC Modeling and System Level Analysis

The material innovation and modeling of super-lattice TECs has been the subject of much recent research [21, 23]. The materials have been thoroughly investigated, focusing on improving the figure of merit ZT and introducing new processing techniques of fabricating TECs [22]. Compact modeling has also been presented in previous works. Detailed SPICE models have been presented in [24, 25], but these models only model the thermoelectric material and do not include the contact resistance effects and integration within a full package. In order to accurately model system performance the development of a 3-D distributed element thermal system is required that takes into account the entire package solution as well as the parasitic resistances that the TEC brings about. Limited full-chip analysis of the integration of TECs in a system has been performed. Long et. al. have explored algorithms for optimal placements of TECs in a chip to maximize the cooling efficiency when they are active [26, 27, 28]. Their work develops a steady state model and integrates it within a simulation environment using an Alpha 21364 microprocessor floorplan. The floorplan is analyzed with power benchmarks and an optimal placement of TECs along with pin assignments and TEC current levels are developed, showing an average of 8°C of hotspot temperature decrease. Although this algorithm does a great job at placing TECs optimally, it fails to consider the transient analysis of TECs and chip temperature. In order to fully model system level behavior, we require a transient model that takes capacitance into account. Limited system level studies have been performed to study the applications of TECs in the dynamic thermal management of microprocessors. Chaparro et. al. [29] have performed system level trace driven architectural studies to understand thin-film TEC integration within a microprocessor. They have simulated a package and TEC together and have studied different control method. They perform trace tests in 3 different configurations, focusing on an extreme performance, ED^2-oriented and a TDP-constrained case.

They account for the total power dissipation and the overhead that the TEC introduces. The extreme performance scenario shows the most promise in performance gains. Although this paper does an excellent job at analyzing the TEC power overhead and execution time gains in a real workload, it does not describe the TEC model in detail or consider the transient behavior of the TEC.

2.5 Energy Efficient Autonomous Operation

A major concern for TEC based on-demand cooling is the need for additional energy. The TEC assisted cooling is necessary only during the chip's thermally critical high power modes, and it is normally turned off during nominal power modes. The finite heat flux generated during the nominal power modes is wasted as the heat energy is dumped in the environment. It is intriguing to note that this heat flux flows through the TEM and can be harvested to generate electrical energy by operating the TEM in the Seebeck mode (TE generator). The switching of the TEM to the TEG mode allows a part of the otherwise wasted heat energy to be reclaimed and stored. The stored energy can be used to provide cooling during the intermittent high power modes. The concept of using a TEM to harvest wasted heat energy has been proposed and studied in prior work [30]. The dynamic mode switching of a TEM for energy harvesting and cooling has first been discussed by Yang et. al in [31] and the performance benefits and implications have been discussed by Choday et. al in [32]. A high-level control system and board level implementation for dynamic mode switching of TEM is presented by Parthasarathy et. al [12]. However, a board-level controller only shows the feasibility of dynamic mode switching, and does not provide a fully integrated low-power solution. The fully integrated solution is necessary to reduce the system volume. Moreover, for chip level embedded cooling the response time in the order of milliseconds is necessary, which is difficult to achieve in a board level implementation.

CHAPTER 3

TEC MODELING FRAMEWORK

3.1   Introduction

Localized high power dissipations (i.e. high heat fluxes) within a chip result in localized high temperature regions or hotspots. The hotspot determines the maximum chip temperature and can set the limit on the maximum allowable power, referred to as the thermal design power (TDP), of a chip/package. In multi-core processors the temperature of the hotspot is time-varying due to the time-varying nature of the heat fluxes. Each core can potentially have a hotspot when running a high-power thread, and hence the location of the hotspots at a given time-instant can change depending on the workload. For a given chip and package, if additional localized cooling can be provided when hotspots appear, that can reduce the maximum temperature and/or relax the constraint on total power limit or TDP. This is referred to as on-demand cooling. The conventional cooling methods using a heat sink and air flow aim to reduce average temperature of the chip and hence, are less efficient in controlling hotspots. Reducing hotspot temperature with air cooling requires a very high flow rate leading to high cooling power and reduced system efficiency. The feasibility of on-demand cooling for hotspots using integrated super-lattice thin-film TECs and evaluate its impact on system performance and the power limit is studied within this chapter. The following chapter will discuss the modeling framework that was established to model the Thermo-Electric Device and the full system. The main contribution is to develop a simulation framework to allow for the device as well as system level modeling in a rapid and efficient manner. Additionally, the simulation framework is able to simulate more rapidly than the physics based simulator at only fractional error.

3.2 TEC Compact Model and Validation

Figure 1 shows the geometry of the chip, package and integrated TEC. The chip and its package are mounted on a PCB while the cooling system is placed on the back side of the chip. The Thermal Interface Material (TIM) improves conduction from the Si backside to the heat spreader. The TEC is embedded within the TIM and makes contact with the heat spreader. The TIM can then be applied on both sides of the spreader and attaches to the chip and heat sink.

3.2.1 Thermal Model for Chip and Package

Due to the well known analog relationship between thermal and electrical properties, we can use circuit level models and SPICE to simulate thermal behavior[33] . More specifically resistors model material conductance, capacitors model specific heat, electrical current density represents heat flux, and the voltage at a given node represents the absolute temperature. Figure 2 shows the three-dimensional distributed RC-based thermal model for the system described in Figure 1. The unit cell for each grid is also shown in Figure 2.
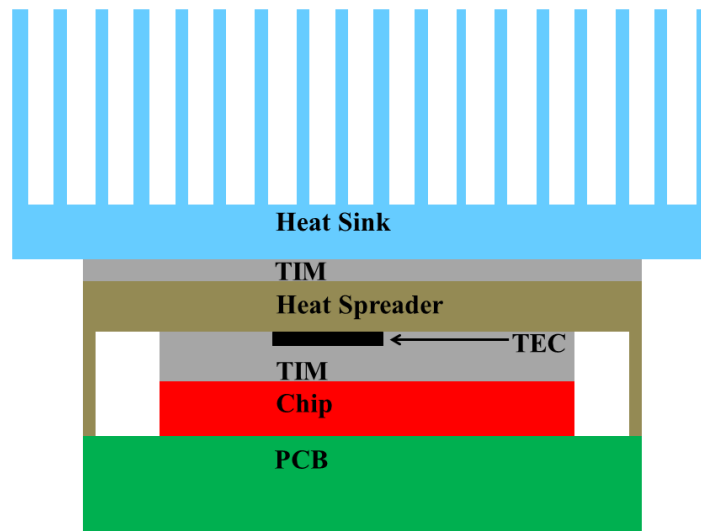
Figure 1: Schematic of a system with a chip, package, and integrated TEC. The TEC is integrated within the Thermal Interface Material (TIM) between the chip and heat spreader.
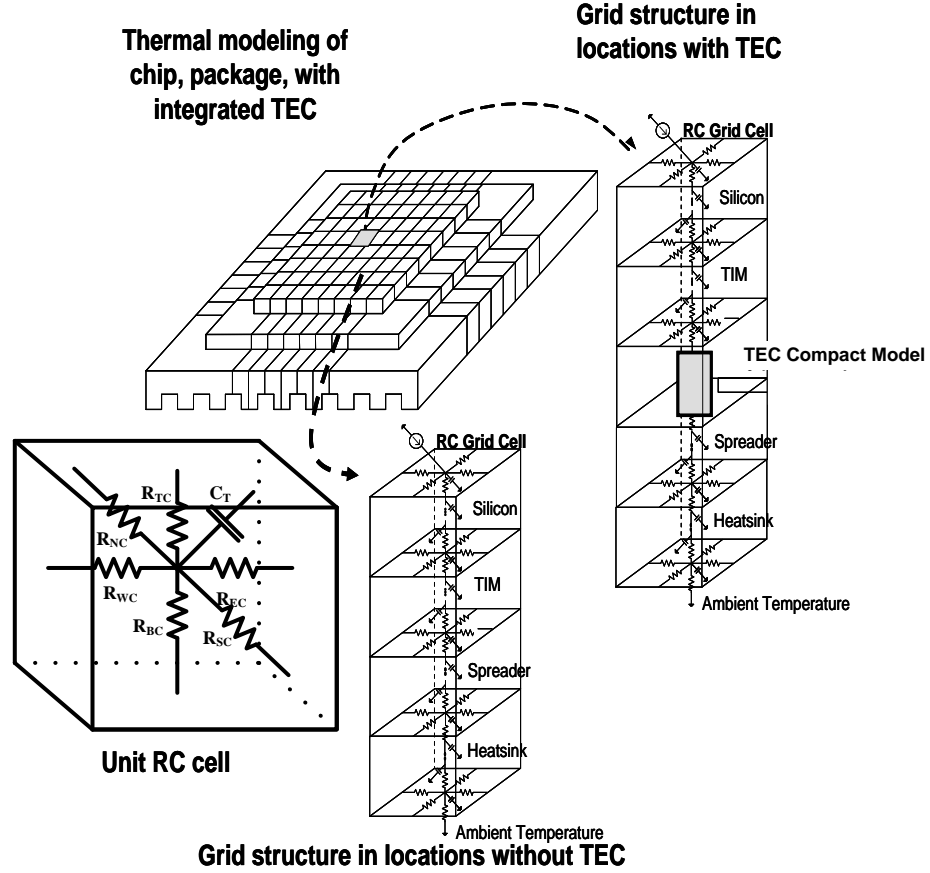
Figure 2: Thermal model for the chip and package showing the unit RC cells, grid structure, and structure of the elements within the TEC location. The model captures both the transient and steady-state system behavior.

The materials are assumed to be isotropic in all directions and the conduction resistances can be calculated based on the grid size and thermal conductivity of the material as shown in equation 1 where L is the thickness of the unit cell and A is the cross sectional area. Capacitance can be calculated from the density of the material, ρ, and the specific heat (c_p) as shown in 1. The material parameters that were used for these calculations are outlined in Table 1.

$$R_{cond} = \frac{L}{\kappa A} \quad and \quad C = \rho c_p \tag{1}$$
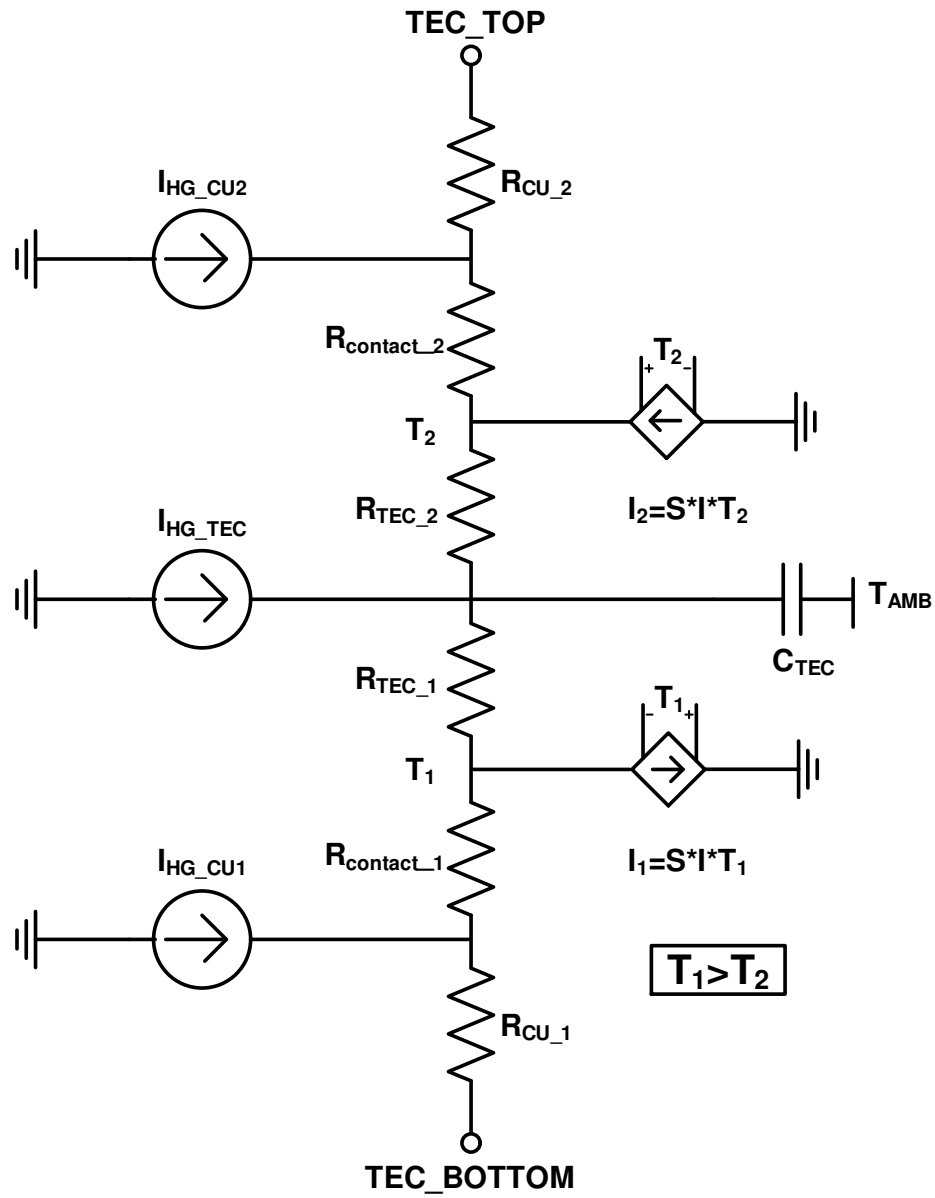
Figure 3: The compact model of the TEC. This model captures both the transient and steady state TEC behavior and considers both the Peltier Cooling mechanics as well as the Joule heating that takes place within the device.

3.2.2 Compact Thermal Model for TEC

With the basic grid for each material defined the only remaining element of our system remaining was the TEC itself. The TEC operates on the basis of the Peltier effect by dumping heat from the hot side to the cold side when a current flows through it. The amount of heat taken away is proportional to the difference in the temperature as well as the current through the TEC. At first glance one would think that by simply increasing the current the TEC would provide large amounts of cooling. Unfortunately that is incorrect as the phenomenon of Joule heating within the TEC brings about I^2R heating that can actually exceed the cooling of the Peltier effect. However, the Peltier effect occurs on shorter time scales than it takes for the Joule heating. [34]

We develop the compact model shown in Figure 3 to capture the above effects. The Peltier cooling of the TEC device was incorporated by adding heat ($\propto SIT_{hot}$) at the hot side and subtracting heat ($\propto SIT_{cold}$) from the cold side of the superlattice, where T_{hot} and T_{cold} are the temperatures of the hot and cold sides, respectively. There are heat generation (HG) sources that model I^2R losses in the TEC as well as the Cu contacts. There are contact resistances representing the parasitic resistances between the Cu and TEC superlattice. The capacitor is lumped and placed in the middle of the TEC. The addition of the capacitor is an improvement on previous works [26, 27, 29] as it allows for the modeling of the transient behavior of the TEC. As illustrated by equation 1 this models the specific heat of the TEC, which defines how much heat is needed to change the temperature.

3.2.3 Model Validation

A full chip with size 9mm X 9mm was created for further investigation. The chip contained a single 3mm X 3mm TEC device in the middle and 3 layers: Si, TIM, and Heat Spreader. The TEC was $100\mu m$ thick and had an $8\mu m$ thick super-lattice material sandwiched between two metallic layers of contacts. The TEC was made of 7x7 n-p couples of p-type Bi_2Te_3/Sb_2Te_3 and n-type $Bi_2Te_3/Bi_2Te_{2.83}Se_{0.17}$. The electrical/thermal contact resistances at the interface of the superlattice-metal layer ($10^{-11}\Omega m^2$; $10^{-6}m^2K/W$) and at

Table 1: System Level Material Properties

3D Distributed RC System Size and Material Properties				
Component Layer	Thickness	Area	Thermal Conductivity [W/mK]	Volumetric Specific Heat Capacity [J/m³K]
Silicon Chip	500 μm	9mm X 9mm	140	1.65M
TIM	125 μm	9mm X 9mm	1.75	1.62M
TEC	100 μm	3mm X 3mm	1.2	0.12M
Heat Spreader	1000 μm	23mm X 23mm	400	3.42M

the interface of the TEC device heat spreader layer ($10^{-10}\Omega m^2$; $8x10^{-6}m^2K/W$) were taken from [21]. The value of S was taken to be $300\mu V/K$ based on experimental measurements in [21]. The heat sink was modeled using a constant convection of 13000 W/m²K and hence an equivalent resistance to the ambient temperature. The input heat flux of the chip was input with a current source at the bottom of the Si layer allowing for ease of changing the total power of the chip. Equations for efficiency of straight rectangular fins [35] modeled the extra spreading that occurred near the edges where the heat spreader (23mm X 23mm) extended over the chip.

We first validated the 1-D TEC model against a Finite Volume Element analysis using FLUENT [36]. The FVM analysis in [36] has been validated against measured results [21]. The TIM, Si, and Heat Spreader were modeled using an equivalent convection resistance and capacitance. Both transient and steady state validations were performed with a relative error of about 2%. Figure 4a shows the transient results and the temperature profile as the TEC is turned on from steady state. We observe that the SPICE and FLUENT curves track very well with a slight underestimation of the temperature in the SPICE model. Figure 4b illustrates that the relative error is kept within 2%. After validating the TEC model, we simulated the compact model for the integrated system considering chip, package, and integrated TEC. We consider the transient temperature of silicon below the TEC for validation. The results were verified with the FVM and the results matched fairly closely as shown in

Figure 5.

3.3 TEC Assisted Cooling

We next analyze the impact of the TEC considering the integrated thermal system (i.e. chip, TIM, heat spreader, heat sink, and TEC). At the system level we concentrate on both steady state and transient power events. Our goal is to understand whether the TEC allows a chip to sustain higher power for a longer time period given a constant temperature target. We consider different cases of system operating conditions when the TEC can increase the system performance and attempt to quantify the performance benefits gained.

3.3.1 Impact of TEC on Steady-State Power: "DC Cooling"

We first study the feasibility of DC cooling. Figure 6 shows the steady-state silicon temperature at the location beneath the TEC as the total chip power is increasing and the TEC is always turned ON. For a target chip temperature, we observe that increasing the TEC current initially allows more power to be dissipated in the chip. For instance, for a target operating temperature of 60 °C, a 3A current through the TEC allows ≈12% additional chip
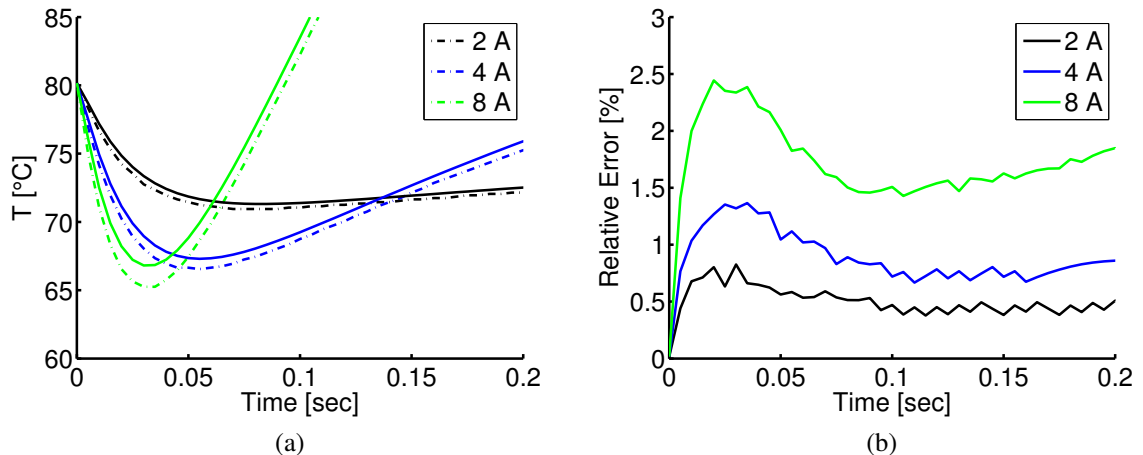
Figure 4: Validation of the TEC model in isolation and with integrated chip/package model: (a) TEC transient behavior solid curves represent FLUENT Finite Volume Analysis and dashed curves are the SPICE model, (b) relative error versus time in the TEC compact model.

16
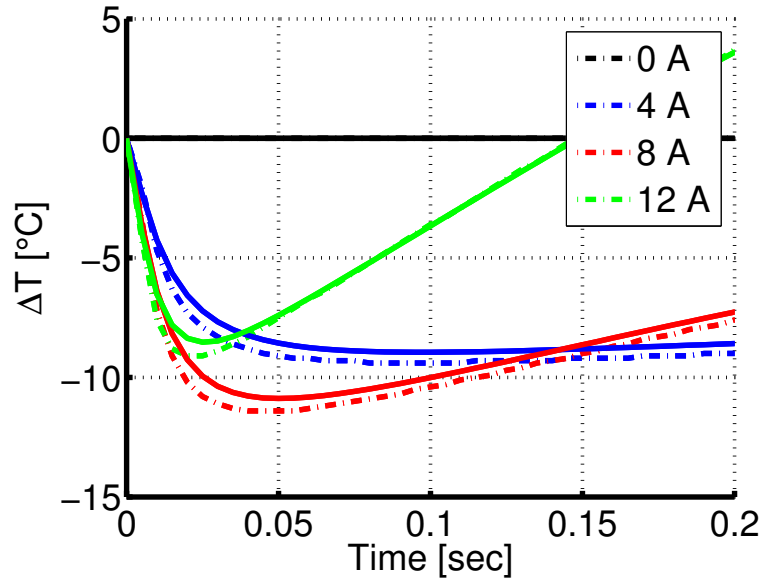
Figure 5: Accuracy of the transient behavior of in the integrated chip, package, and TEC mode of the chip/package model.
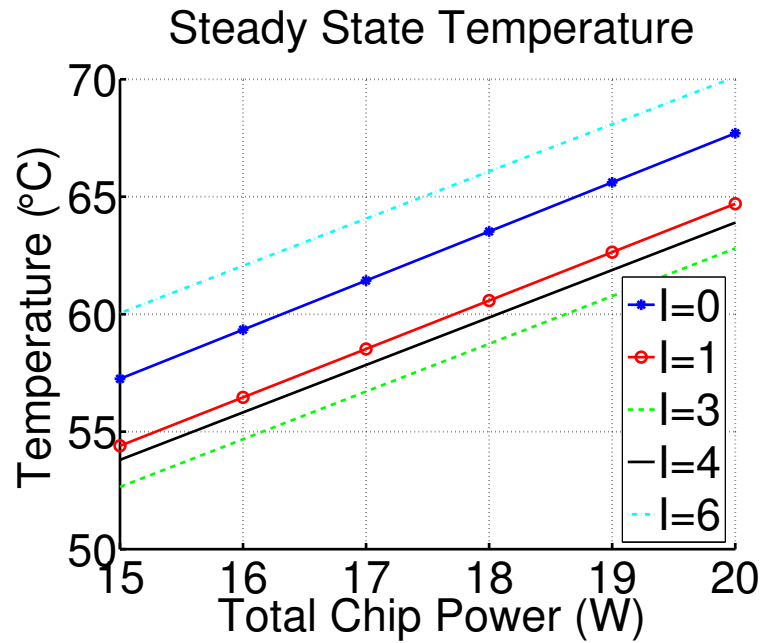
Figure 6: Steady State Temperature directly under the TEC.
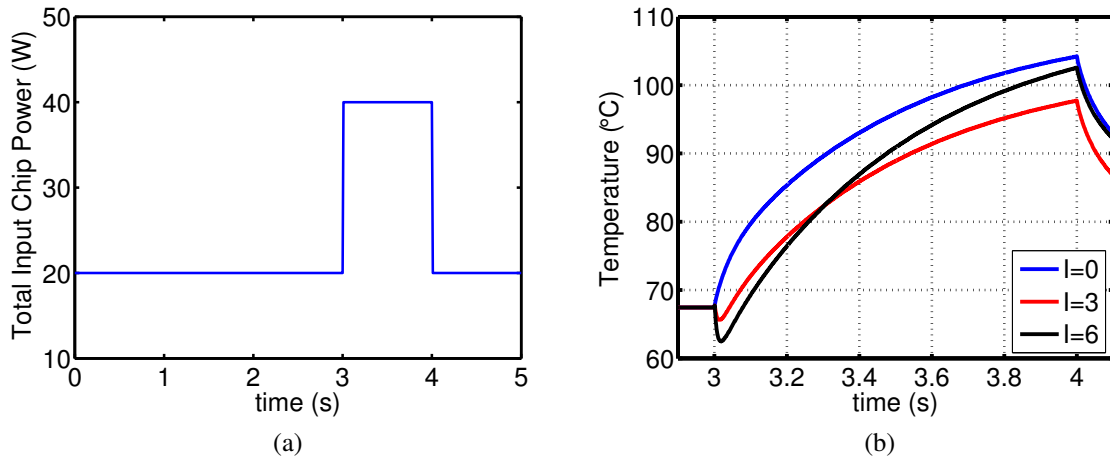
Figure 7: Power pulse applied to the chip modeling a heavy workload (a) the nature of the power pulse and (b) the transient response of chip and package with a 1 second power pulse applied.

power compared to having no current (I=0) in the TEC. The extra power can be exploited to increase steady-state performance (e.g. higher clock frequency) for a given package. This can significantly raise the TDP of the processor and allow a higher baseline performance. Therefore if the overall system has the power budget, the TEC can always be operated with the optimal current through the device. The designer needs to take note that further increase in the TEC current introduces high Joule heating and the chip reaches a higher steady state temperature for same power. As we can see, above 6A of current into the TEC, the Joule heating is actually greater than the Peltier cooling and the chip actually has a worse baseline temperature. The problem with 'DC' cooling is that the TEC will always be dissipating additional energy, which might decrease the overall system performance when the chip is not operating under thermal stress or is in idle mode. This means that the TEC should be used only when the chip is operating near the thermal limit.

3.3.2 Impact of the TEC on System Level Power/Performance: "Transient Cooling"

As mentioned, from a system design perspective it does not make sense to have the TEC turned on at all times. With this in mind, we next study the implications of TEC assisted cooling when the chip receives a large power spike (Figure 7). It is important to note

that even adding a TEC without any current flowing through it (i.e., I=0A) is beneficial compared to the no TEC (package only in Figure 8) case, as the TECs copper contacts provide better heat conduction than the TIM material (Fig 7b, simulations with HTC=2500 W/m^{2}K) and allow for a lower base line temperature for a given chip power. With a finite TEC current (I=3A), we can see an initial dip in the temperature when the TEC turns on. After a time interval the temperature starts rising again and ultimately reaches a lower steady state temperature compared to the I=0A case. The initial reduction in temperature is due to the fact that the Peltier cooling starts operating immediately after the TEC is turned on, while the response time for the Joule heating is marginally higher. When the Joule heating becomes significant, the temperature starts rising at a faster rate. As the current is increased further (3A to 6A) we see a greater initial cooling, but the rate of increase after the initial dip is much higher and hence the temperature becomes higher than the 3A case. The transient cooling obtained with the TEC can significantly reduce the rate of increase of the chip temperature and hence, can sustain a power pulse for longer duration without violating thermal constraints. For example, Figure 7b shows that a TEC current of 6A can allow the system to remain under a target temperature of 75 °C for an extended time (Δt in the order of 100ms) period compared to the no TEC case; 100's of ms of time-extension is quite significant at the processor time scales.

3.4   Principles of Transient Control

The previous section has briefly introduced the potential applications that TECs possess in increasing the TDP of a package and delivering higher power dissipation in the package. The transient effects of the TEC were presented and their potential briefly introduced. However, there are many challenges in efficient control of the TECs transient behavior, as there is direct dependence with the workload from a processor as well the architectural design of the chip. This introduced many cases that the TEC might be able to increase the performance of a processor or workload. In this section we discuss the different challenges

19
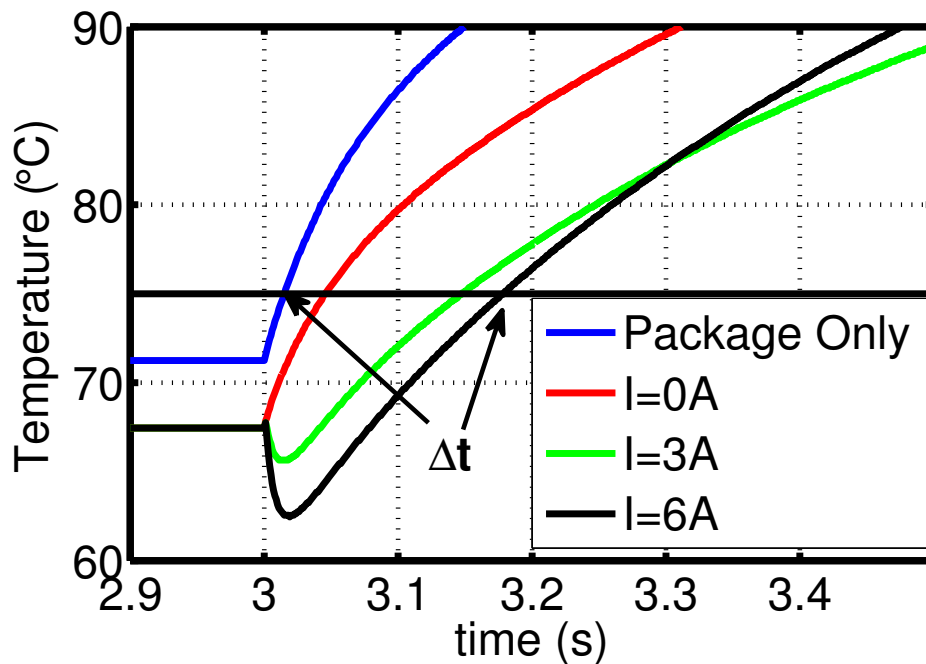
Figure 8: Transient Response and the concept of workload extension.

and principles for efficient control of TECs to exploit their potential in managing transient temperature variations. Due to the enormous design space and architectural intricacies of systems we focus our efforts on a set of specific cases. First we consider that the time duration of a power pulse is known and our objective is to optimally choose a TEC current level to minimize the temperature for that power pulse. This is a case that typically occurs in a high performance processor in the so called Turbo-Boosting that Intel has developed. [1] The main concept of turbo boosting is to take advantage of the transient behavior of the microprocessor package, and allow the system to operate beyond the TDP limit for a short period of time. This takes advantage of the thermal capacitance of the package and the delayed RC response. The performance benefits are realized through dynamic frequency and voltage scaling (DVFS) by increasing the frequency and voltage of the processor to stay under thermal budget. Figure 9 shows this concept graphically. As we can see, when the core temperature is low, the processor is put into a Turbo state, where it dissipates about 20-30% higher power than the allowed TDP budget. The core temperature begins to

increase but takes some time to actually reach the TDP limit, due to the thermal capaci-
tance of the package. Once the processor starts getting closer to the threshold, the power
is incrementally scaled back to its TDP limit and the processor is allowed to remain under
the TDP limit. This allows a performance gain while staying under the thermal budget. It
is also important to understand the performance gains that are associated with using turbo
boosting. Due to the complex architectural design, instruction set scheduling, and memory
hierarchies, scaling the frequency of a core will not result in a linear performance gain with
the frequency. Extensive work has been done in the architectural community to assess the
performance gains of the turbo boost feature. Charles et.al. [37] have through an exten-
sive analysis characterized the turbo-boost behavior in various workloads. In a single core
workload, they show that on average, CPU intensive workloads achieve a speed up of about
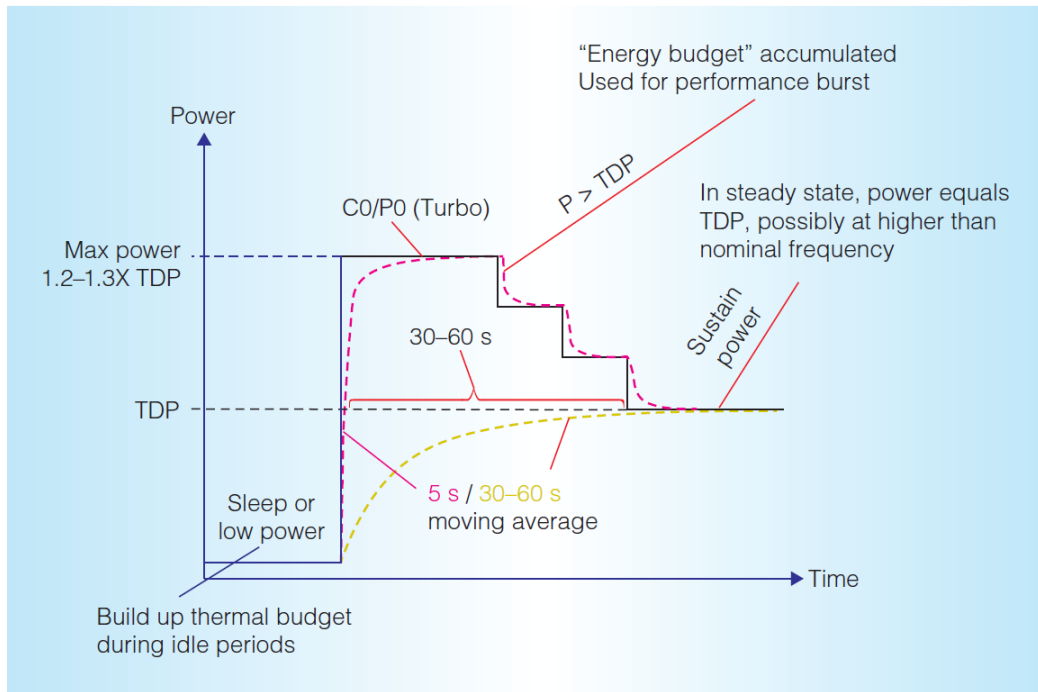
Figure 9: Transient behavior of the Intel Turbo-Boost technology. If the system is at low
temperature, the TDP can be exceeded to achieve a very high performance and take advan-
tage of the thermal time constant of the package. Once the temperature reaches the TDP
limit, the processor is put back into the nominal DVFS setting. Adapted from [1]

21

6.11%. This gain decreases to 5.15% in moderate memory-intensive loads and further reduces to 2.93% in highly memory intensive workloads. These results are achieved from a 16% increase in total system energy dissipation. These results serve as a proxy for us to assess the performance gain benefits of the TEC.

3.4.1 Single-Core Turbo Boosting

In this section we evaluate the performance level benefit of using the TEC in a package and its effect on the turbo boosting time. We first study the scenario that the processor needs to be in the turbo boost mode in a high-power mode for a target time-period. Our objective is to minimize the temperature during that period. The TEC controller here turns the TEC on when the power pulse arrives and turns it off after the power pulse disappears. From an architectural perspective, this implies that once the workload is scheduled to a core and there is knowledge of a high power event, the TEC is also turned on. When the power pulse is over (the workload has completed), the TEC is turned off to minimize Joule heating and power dissipation in the TEC. The goal here is to achieve minimum temperature at the end of a power pulse. For our experiments we varied the pulse width from 100ms to 900ms. We noted the temperature difference at the end of the pulse between the no current in the TEC case (I=0, baseline) and different values of TEC current up to 15A. The highest temperature difference value was taken to be the optimum TEC current for the given target pulse width. Figure 10b shows the required optimal TEC current to achieve the minimum temperature at the end of the pulse width and Figure 10c shows that maximum achievable cooling (compared to the I=0) case. We observe that as the pulse width increases the optimum level of current also decreases. This is because the higher currents through the TEC cannot sustain their initial cooling levels for long periods of time. This is because the Peltier effect provides initial cooling rather quickly, but the Joule heating within the TEC eventually begins to reduce these cooling effects. Note that with the increasing pulse width, the optimal current approaches the steady state optimum condition (i.e. I=3A as shown in Figure 6). We further note that the maximum achievable cooling
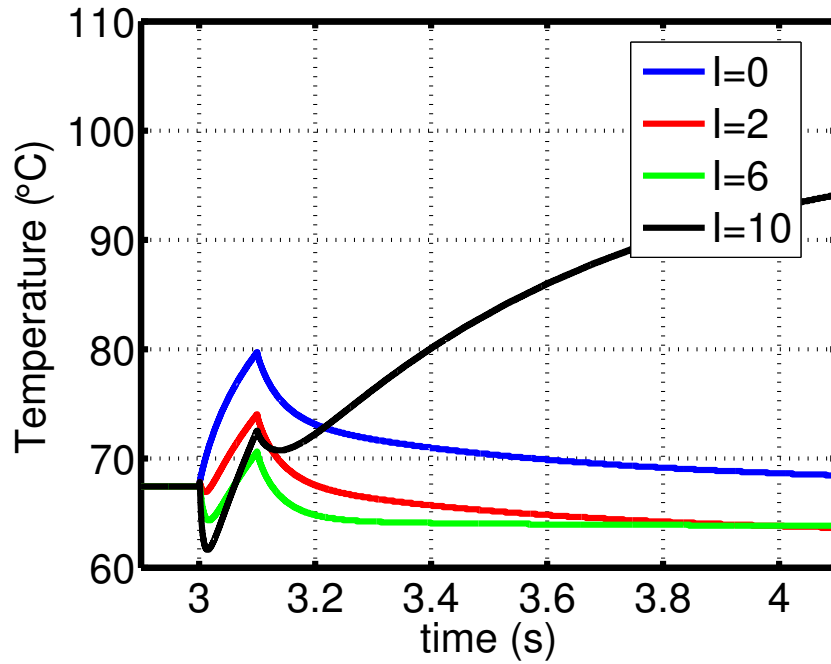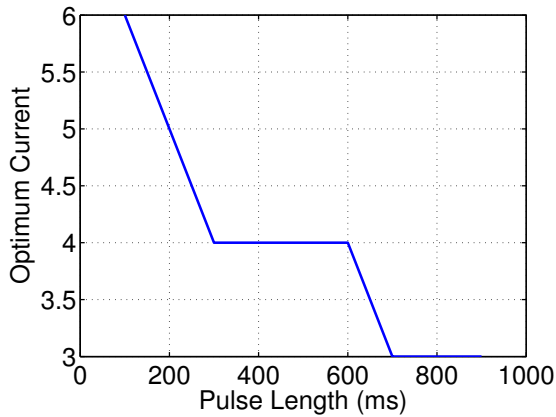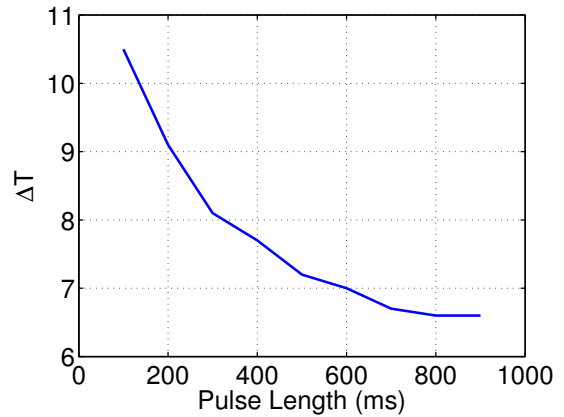
Figure 10: Principles for TEC control for high power pulses of known width: (a) transient response of chip and package with a 100ms pulse applied to the chip showing the need for turning off the TEC with the power pulse (Joule heating), (b) optimum current for minimum temperature (i.e maximum cooling compared to no TEC current case) at the end of a pulse, and (c) maximum cooling at the end of the pulse (i.e. ΔT with respect to no TEC current) with varying pulse width.
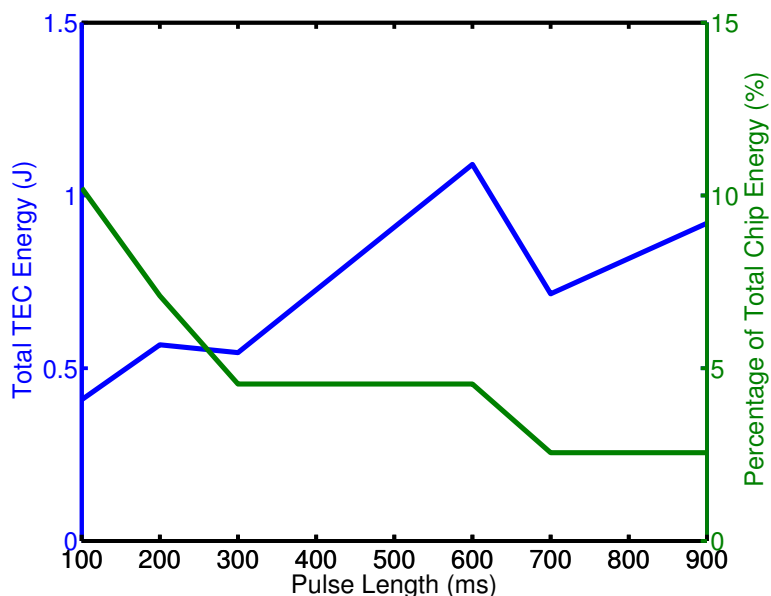
Figure 11: The TEC energy dissipation considering the current flowing through the TEC and contact resistances. The blue curve corresponds to the actual TEC energy (left axis) while the green line corresponds to the percentage of the TEC energy to chip energy dissipation for that pulse duration

also reduces at higher pulse width. Figure 11 shows the TEC energy dissipated (due to the TEC current flowing through the finite resistance of the TECs and the electrical contacts) at different pulse width using the optimal TEC current. We observe that the additional TEC energy dissipation is within 10% of the energy dissipated by the power pulse (40W pulse) within the pulse duration interval. Therefore, we conclude that with the optimal level of TEC current it is possible to minimize the chip temperature for different durations of power pulses. This allows for a look-up table (LUT) based solution where the optimal TEC current to be pulsed is determined based on the workload to be scheduled, and the solution can be implemented directly at the architecture level.

3.4.2 Multi-Core Turbo Boosting

Modern high performance microprocessors have multiple cores and the number of cores continue to increase within new generations. This creates additional challenges in managing the thermal transients on chip and creating an effective cooling solution. This section
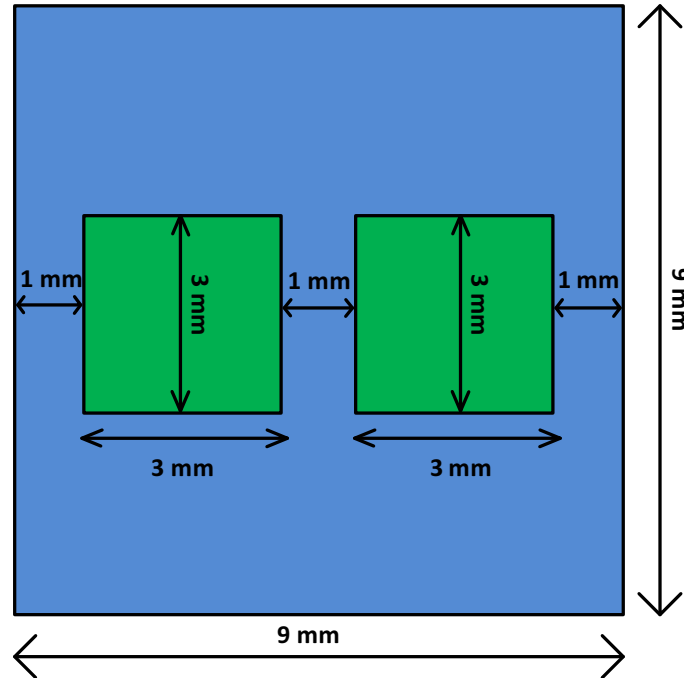
Figure 12: The multi-core simualtion framework. A 9mmX9mm chip is used with two 3mmX3mm cores placed centrally. Each of the cores has a TEC completely covering it.

will show how TECs can help mitigate multi-core thermal challenges and attempt to quantify performance gains in a multi-core scenario. The first step in modeling a multi-core is to create a simulation framework. Figure 12 shows the chip level distributed 3D RC grid. We have a 9mm X 9mm full chip with the same heat spreader previously used and we place 2 3mm X 3mm cores centrally. Each core is also covered fully by the 3mm X 3mm TEC used thus far. This models a high performance 2 core system. The first step in the analysis is to quantify the thermal challenges that arise from having multiple cores on chip. For this simulation, we assume that the chip is operating with 10W of total background leakage power, and apply a maximum load to each core of 15W. This power value is roughly the power that a single core dissipates in the latest 4th generation Haswell line of processors from Intel. Because the cores are in close proximity and the Silicon layer has very high thermal conductivity, the cores will not only heat their local area, but a core executing a high power workload will heat up a neighboring core, thereby increasing its temperature. We can see this effect in figure 13. We show the power profile that is applied

at each core and the resulting transient temperature. We turn on core 1 to full power and observe the transient As we can see initially the first core is turned on and starts to heat up, while also heating up the second core which is curently sitting idle. After some time, the cores reach their steady state value, and both temperatures are below the required junction temperature limit, typically around $100 - 110\,°C$. After the second core is turned on, we observe a similar case and the second core starts to heat up the first core. Both temperatures eventually reach a steady state which is near the junction limit. This means that in a dual core scenario, both cores will not be able to be operated with a turbo-boost feature, as they would exceed the junction temperature limit. This can be a severe performance limiter. In order to combat the thermal limits in a multi-core environment, computer architects have began to use dynamic thermal management technique called core hopping. This means that anytime there is an idle core not being used, and the current core where the workload is being executed is nearing it's thermal limit, the workload can be directly mapped to the idle low temperature core. This will allow the current core to cool down, and balance the average temperature across the chip. This in effect is a way of balancing the power on the chip. Of course it is clear that this technique will have overhead that needs to be accounted for. In order to migrate a thread to a different core the data needs to be transferred and the registers need to be properly written, the shared caches need to be updated and trained, and the branch predictors need to be trained [38]. If the second core is "cold" and there are many branch mis-predictions and cache misses, the overhead can also significantly increase. Donald et.al. [39] modeled core migrations in a modern 4-core 3.6GHz 90nm design and have used a migration time of 100μs. This is not very significant at thermal scales and for added worst-case scenario, our simulations will assume a 1ms core migration time. This allows for any additional overheads that might be associated with the temperature sensing on-chip and additional hardware controllers. To assess the TEC effectiveness in a multi-core scenario, we revert back to the simulation framework in figure 12 and incorporate the architectural considerations. We assume a workload of very
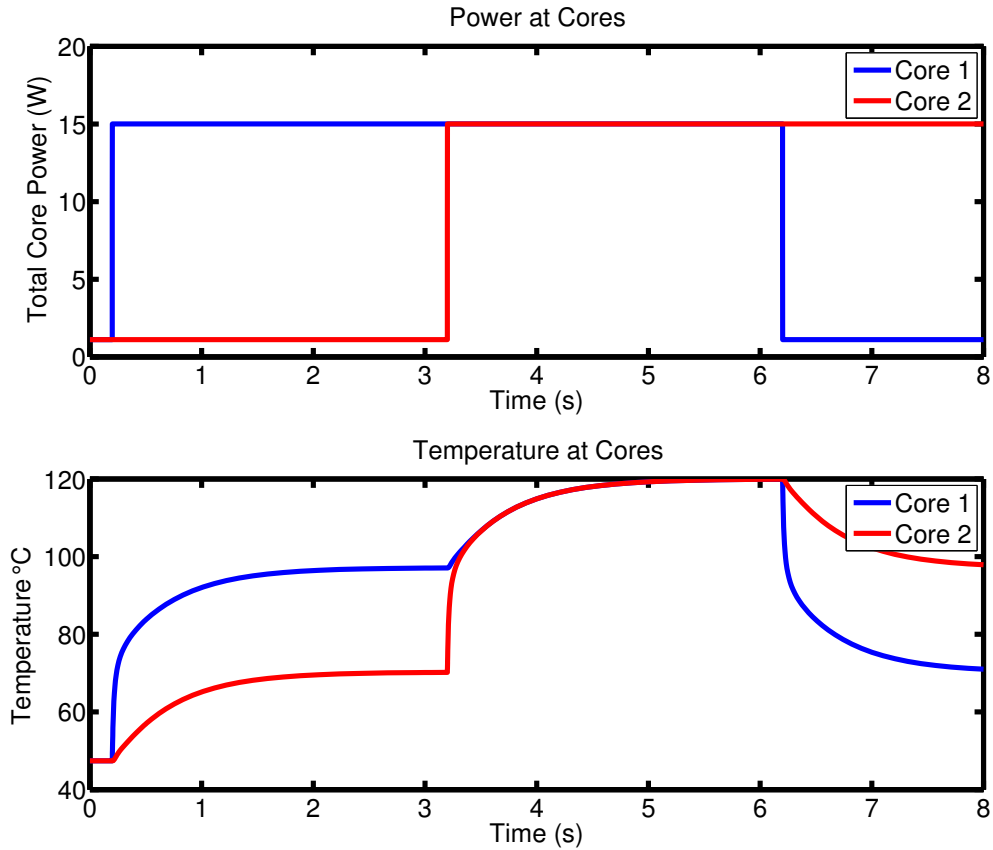
Figure 13: The multi-core transient simulation. Each core has 15W of power applied. We observe that when only core1 is on it's temperature rises to about $100\,°C$ while core 2 heats up about $20\,°C$. When both cores are on, their temperature reaches $120\,°C$ and exceeds the thermal limit.

long length, which models a heavy computational workload such a scientific computing. The goal is to complete the workload with the highest performance, thereby the core needs to be operating a the maximum frequency and voltage and hence maximum power. As we saw previously, with both cores running at 15W of power, the temperature is exceeded and this causes a thermal violation. We therefore need to balance the thermals as previously discussed by migrating between the cores. Figure 14 shows the power profile at each core. The TEC is also pulsed with the same current profile, in order to provide cooling to the core that is executing the workload. We use a fixed current through the TEC. In order to achieve highest performance, we seek to maximize the power at each core while keeping
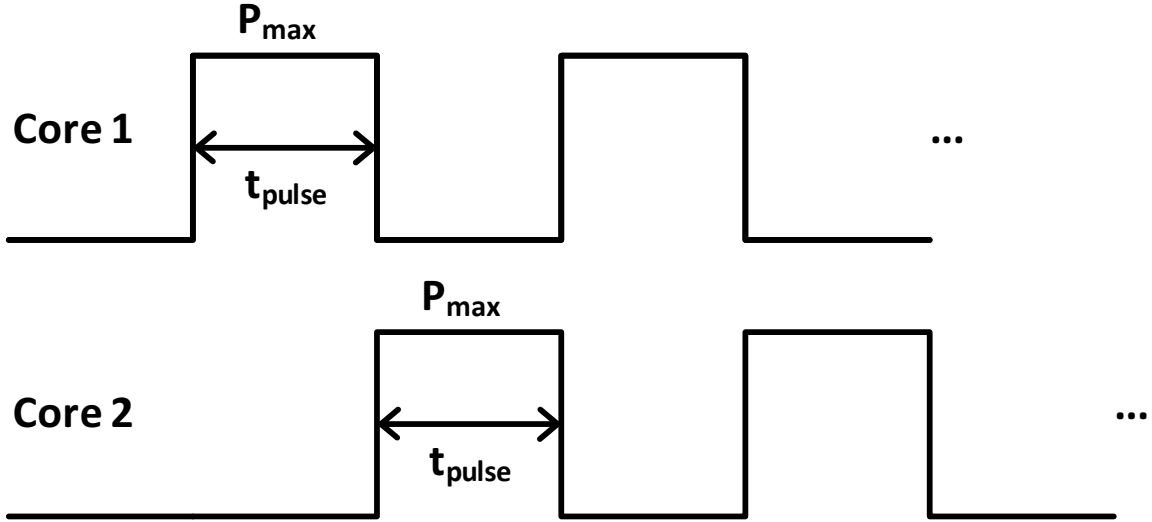
Figure 14: The multi-core transient power profile. Each core has maximum power pulsed, after which the power is switched to the other core. Each TEC is also turned on in the same fashion with a fixed current.

the temperature below the thermal limit. Figure 15 shows the transient simulation result of the multi-core scenario with P_{MAX} = 15W and t_{pulse} = 200ms. As we can see by switching the workload from core to core, we never let the core reach it's steady state temperature limit and instead manage the temperature within bounds at both cores. In this case the maximum temperature that each core reaches is 82 °C, while without migration the single core reached 100 °C. Since the thread migration of 1ms is only a 0.5% overhead this presents a very low overhead to reduce the temperature significantly.

The next interesting aspect to consider is how the pulse length and more importantly TEC current can lead to a reduction in the chip temperature. In order to be consistent across simulations, we define T_{max} as the maximum temperature that each core reaches, and due to the symmetry of operation each core reaches the same maximum temperature. We study the effect of the power pulse height, length and the TEC current on T_{max}. Figure 16 shows the effect of varying all the aforementioned effects. We observe that without a TEC in the system, the cores are operating very close to the junction temperature limit and require the power pulse width to be very small. When we insert a TEC into the system we see an average temperature reduction of about 15 °C as compared to the package only case. We

also observe that the TEC current has a significant impact on the maximum temperature of the cores. As we can see in the figure, increasing the TEC current from 1A to 4A has a net benefit as it reduces the temperature across all power pulses. Increasing the current beyond 6A actually begins to increase T_{max} when compared to the 4A case. This is due to the Joule heating effects of the TEC as after a very long time of pulsing the TEC current this has reached semi-steady state. This effect can be observed by looking back at figure 8, where the cooling level of the initial TEC pulses provide a more pronounced cooling than the later ones. Since the workload is of very long length, we can neglect the initial transient
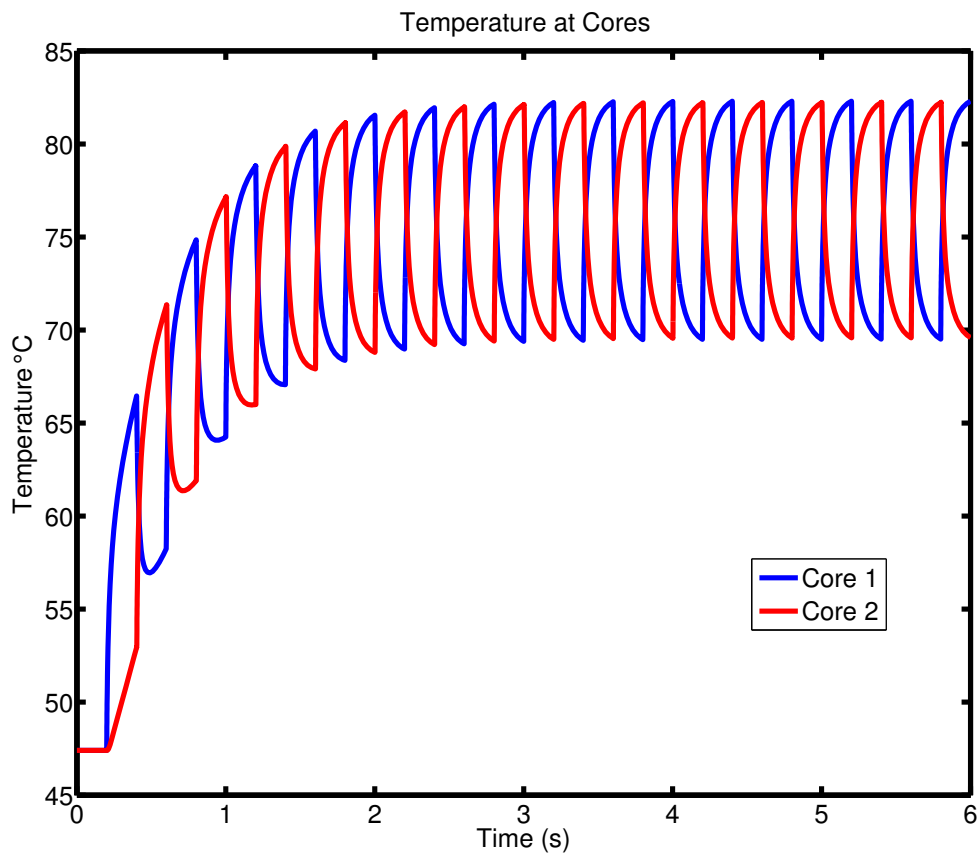
Figure 15: The multi-core transient temperature response. Each core has maximum power pulsed, after which the power is switched to the other core. The temperature at each core varies by about 15 °C and the each core reaches the same maximum temperature. By switching the power between cores, we also prevent each core from reaching its higher steady state temperature.
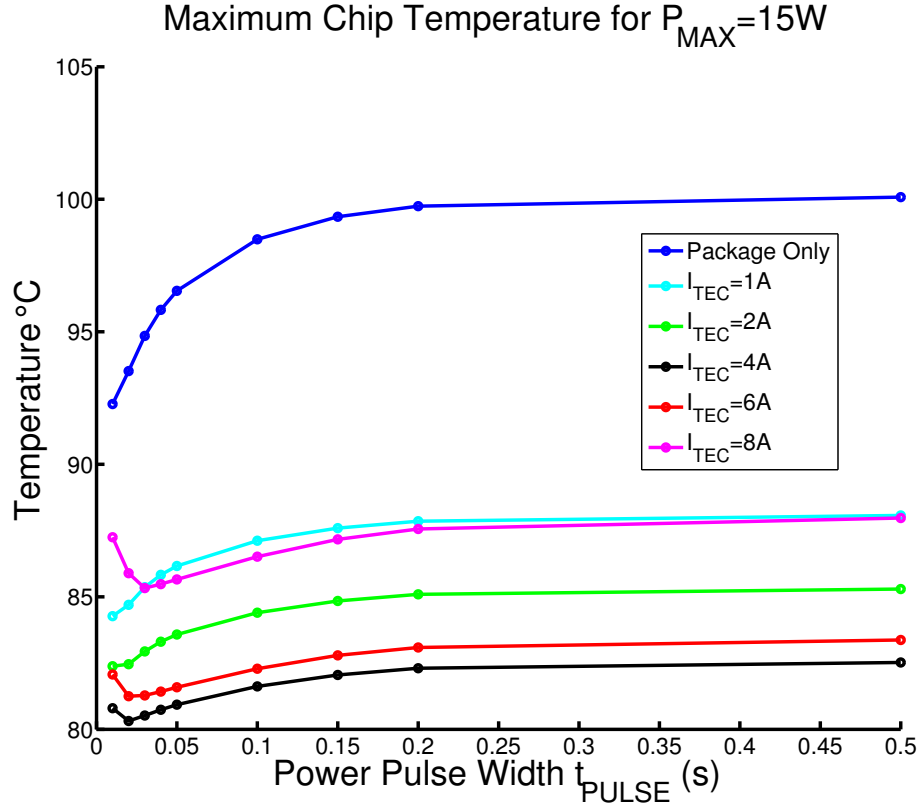
Figure 16: The multi-core framework with 15W applied at the cores. The effect of the TEC current level on the overall maximum temperature T_{MAX} is shown.

and focus on the temperature profile beyond 4s. We can define an optimum current for a given power across the cores that will lead to the minimum temperature of each core. In this package with the TEC size and characteristics, the optimum TEC current $I_{OPT} = 4.5A$. Similarly, for any package and expected power patterns a system designer could use this framework to find an optimum TEC current.

It is clear that using the TEC in the system allows the silicon to have higher power dissipation and hence higher performance. Figure 17 shows the effect of core power on T_{max}. We have used the optimum TEC current level for these simulations. We can see that by using the TEC in the system we can allow each core to achieve higher power and increase the performance while staying under the thermal limit. Additionally, a core can gain an additional performance boost by reducing t_{pulse} to its optimal point. Take for example the case of $P_{MAX} = 17.1W$, where with a long pulse width of 0.5s, we can keep the temperature
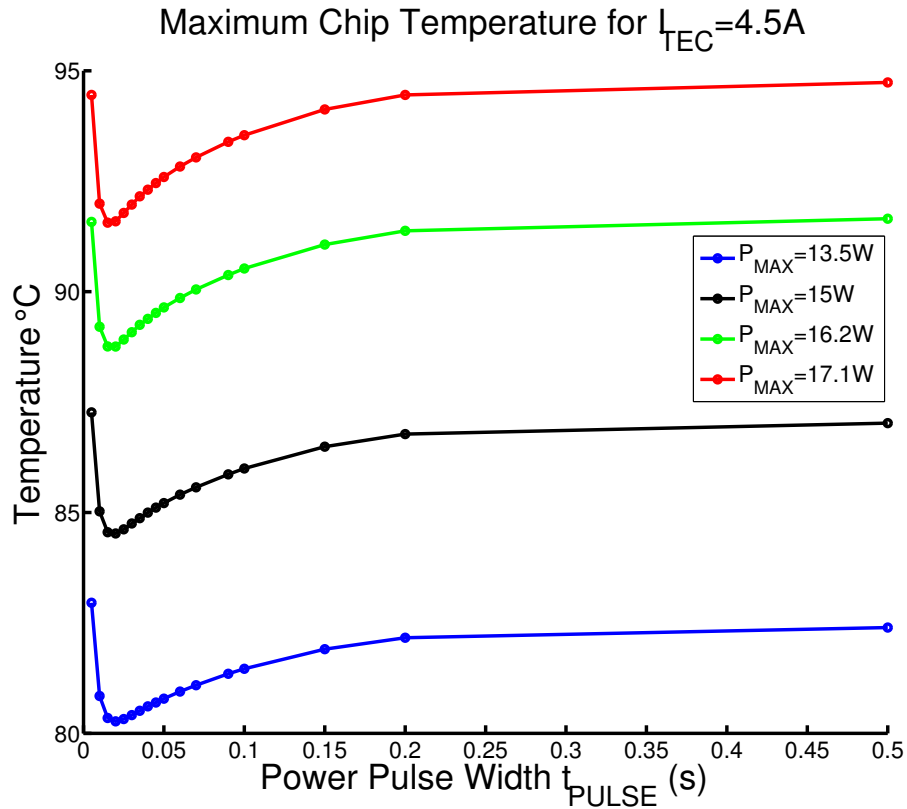
Figure 17: The effect of increasing the power at each core. We see that the TEC allows us to dissipate higher power in each core, while staying below the temperature limit. Additionaly we can sustain an even higher performance by operating the cores at the optimim t_{pulse} as there is further reduction in temperature.

of the chip right below 95 °C, but decreasing the pulse width to 20ms can yield an additional decrease in the chip temperature by 3 °C. This means that if we used the optimum pulse width we could sustain an even higher power chip power, thereby increasing performance. The trade-off here is that the overhead of the migration time will be greater as we are incur a 1ms penalty with every migration. If the workload is hopping from core to core every 20ms, this is a 5% overhead. Since the overall performance gain is strongly dependent on the chip architecture, we cannot choose the most optimal design point. Rather by using this framework, an architect can study the TEC effects on a particular microarchitecture and choose an optimum design point.

3.5 Summary

In this chapter, we have fully presented a modeling framework for the evaluation of a TEC. We have studied the prospect of on-demand cooling with super-lattice thin-film TECs integrated in a chip and package. We have developed compact models of TECs and integrated that in a full-chip 3D thermal RC model. The integrated structure is used to perform steady-state and transient thermal simulations considering time-varying power pulses. We have observed that using a TEC for steady-state cooling has marginal impact, while transient cooling allows a processor to sustain a high power pulse for a longer period of time without violation of the thermal threshold. This can significantly reduce the thermal events in processors and help improve thermally limited system performance. We have presented a methodology to evaluate how a TEC embedded in a system can help alleviate thermal challenges. In the single-core scenario, we present a methodology to evaluate the optimum current for a single known length power event. This is an event that is coupled with the architecture of the system. We showed that for an event of any length, the TEC provides a net reduction in temperature and allows for higher power dissipation. In the multi-core scenario, we showed that using a TEC in the system can not only help sustain a higher power in each core, but that many trade-offs exists in the allocation of the power to allow for minimum temperature at each core.

CHAPTER 4

ON-DEMAND TEC ASSISTED COOLING AND CONTROLLER DESIGN

4.1 Introduction

In this section we present the principles for efficient control of TECs to exploit their potential in managing transient temperature. This is a control problem of reducing the thermal events (i.e. sustain a power pulse for a longer period of time) before the chip temperature crosses a threshold and the need to invoke Dynamic Thermal Management (DTM) arises. The control principles here are based on sensing the temperature to activate or deactivate the TECs. We restrict the study to binary control i.e. TEC is either off or on with a fixed current. Note that turning the TEC off helps reduce the energy dissipated in the TEC. We consider that the fast sampling rate and high resolution of the sensors offer the possibility of very fine control of the activation and deactivation of the TECs. This is justified as advanced on-chip temperature sensors have sub-millisecond sampling time and less than $0.5°C$ resolution [40, 41]. We first consider a Threshold Based Controller (ThBC) where sensed temperature is compared against a threshold. The TEC is turned on when the temperature crosses a certain threshold and it is turned off when temperature is less than the threshold. The controller keeps the temperature near the threshold for a significant period of time and delays the occurrence of the thermal event. However, the TEC suffers from continuous on/off transitions (limited by the sampling frequency of the sensor) that could lead to degraded TEC reliability. We explore an alternative approach, referred to as the Maximum Cooling Based Controller (MCBC), to minimize the number of transitions. Here we turn the TEC on beyond a threshold temperature but allow the temperature to reach its minimum value before turning the TEC off. This helps exploit the initial high transient cooling offered by the TECs. The Joule heating starts causing additional heating past the minimum point and turning the TEC off is more efficient. The control principle
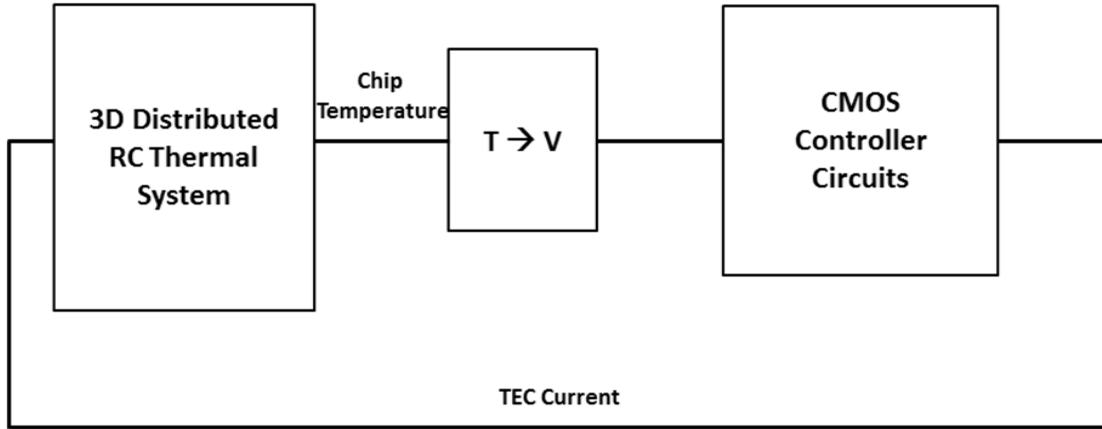
Figure 18: Co-Simulation Framework with integrated CMOS controller.

is to turn the TEC on when the threshold temperature is reached, turn it off after the minima is reached, and turn in back on when temperature crosses the threshold again. The circuit level designs and simulations of the ThBC and MCBC controllers are performed in 130nm CMOS technology. After verifying the controllers electrical characteristics in isolation, they are integrated with the distributed RC based full package thermal models to perform electro-thermal co-simulation. Figure 18 describes the co-simulation framework. The electro-thermal co-simulation helps accurately characterize the effect of the ThBC and MCBC based TEC control on transient temperature considering the electrical response time of the sensors and controllers as well as energy overhead of the controllers. The designs of the controllers are addressed in detail in following sections. A power pulse with 10W baseline power and a high power of 42.4W was applied (see Figure 7a). The high power pulse was applied at 100ms and remained on until the end of the simulation. The total heat flux of the power pulse was 52.35W/cm^2. The heat transfer coefficient (HTC) used to model the system was set to 625W/m^{2}K, representing the low end of the cooling capabilities provided by commercial heat sink solutions.
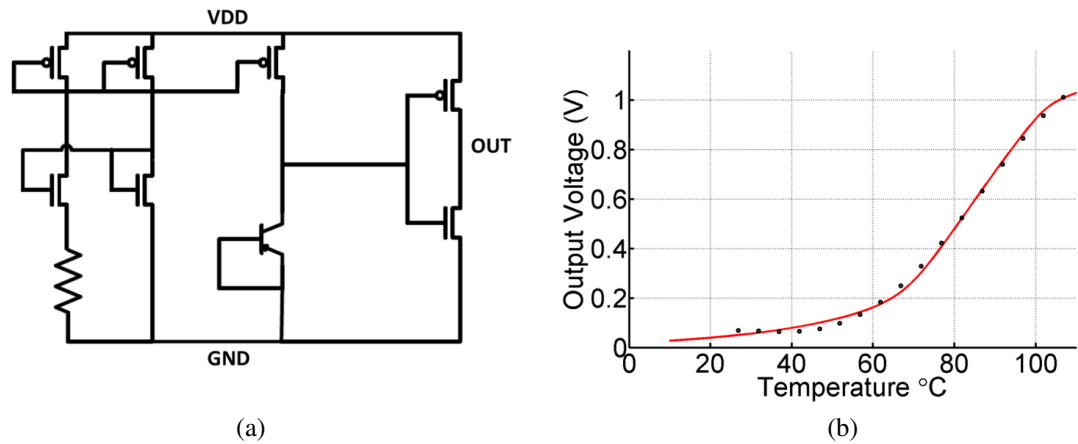
Figure 19: Temperature Sensor Circuit: (a) schematic and (b) output characteristic and polynomial fit.

4.2 Temperature Sensor Design and TEC Electrical Model

A temperature sensor was implemented with a constant current source going into a diode connected BJT and generating a known voltage at a given temperature (Figure 5a). Because the voltage measured across the diode only varies by about 1-2mV/K and we are only interested in a small range of temperatures for controlling the TEC, we have added an amplifier to the output of the diode. This amplifier provides rail to rail output voltage within the high gain region falling between 60 °C and 100 °C. This will ensure that the sensor output will be close to ground when the chip is running in a low power state and close to VDD at very high chip temperatures. The output characteristic of the temperature sensor is shown in Figure . As we can see the sensors high gain region is located between $60 - 100$ °C. We fitted a 4th order polynomial to the sensor output so that the temperature to voltage conversion could be done within our electro-thermal co-simulation environment (in the co-simulation environment the temperature is represented by the voltage at the nodes of the distributed RC grid). The polynomial fit is also shown in figure . The voltage of the Silicon layer from the RC thermal grid is an input for the polynomial source, which in turn converts the voltage between 0 and 1.2V to provide input to the controllers.
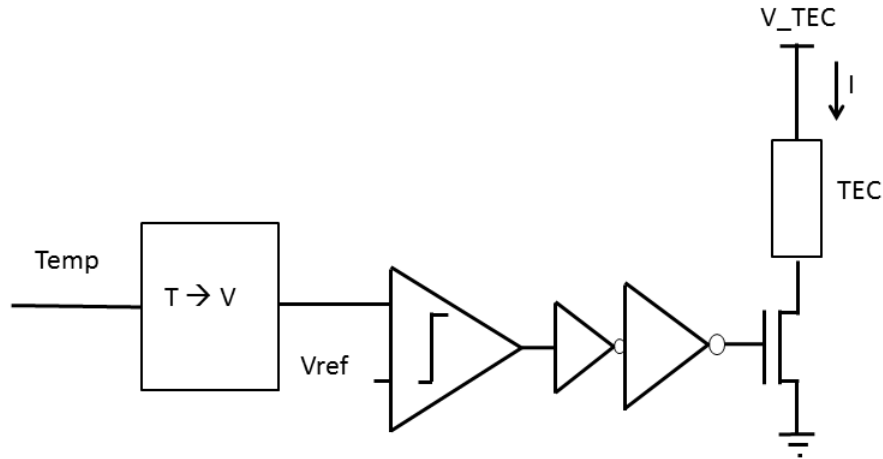
Figure 20: Threshold Based Controller (ThBC) circuit schematic.

4.3 TEC Controller Design

4.3.1 Threshold Based Controller (ThBC)

Figure 20 shows a top level schematic of the ThBC. The temperature to voltage converter is a polynomial fit of a simple BJT based temperature sensor, and it converts the chip temperature to voltage. This voltage is compared to a reference representing a predefined control temperature using a comparator. The current through the TEC will be set by the TEC supply voltage.

The ThBC was integrated with the full package thermal model in order to run full system simulations. The current through the power FET is fed as an input to the thermal TEC compact model and the TEC is turned on when the comparator output is a logic high. As figure 21 shows, we can see regulation of the silicon temperature (simulations with HTC=625 W/m^{2}K). Once the temperature exceeds the predefined threshold, the TEC turns
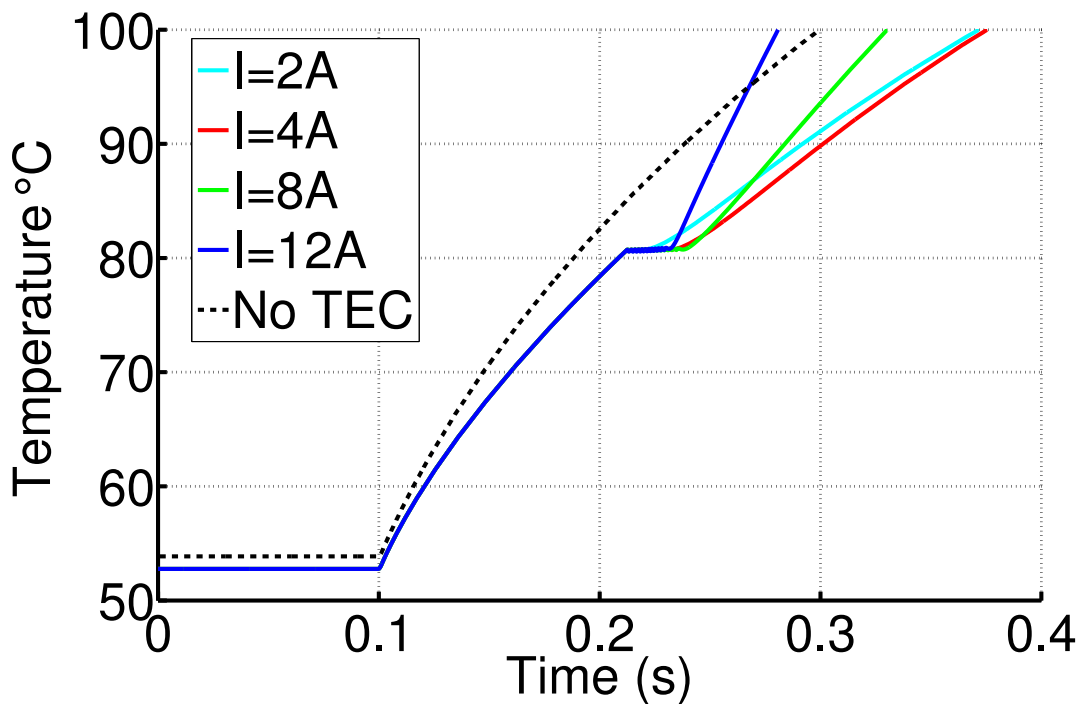
Figure 21: Threshold Based Controller (ThBC) simulation results.

on and immediately starts to decrease the temperature. Once the temperature falls below the threshold, the TEC turns off and the silicon temperature starts increasing again. When the generated heat is high enough that that the TEC cannot reduce the silicon temperature, the TEC remains on but the temperature also continues to increase. Figure 21 shows the temperature response of the integrated system. It can be observed that at low TEC currents, the power pulse can be sustained for shorter times. Increasing the current initially helps to increase the time-duration as shown. However, at higher current the Joule heating contribution is also higher. As we can see the rate of increase due to Joule heating at 12A is much greater than the rate at 8A. Consequently, increasing the current beyond a certain point actually reduces the overall extension period. Therefore, a careful choice of the TEC current is required to maximize the extension time. Note that once the temperature exceeds the threshold of the microprocessor, the DTM or throttling mechanism will be invoked to reduce the power dissipation itself.
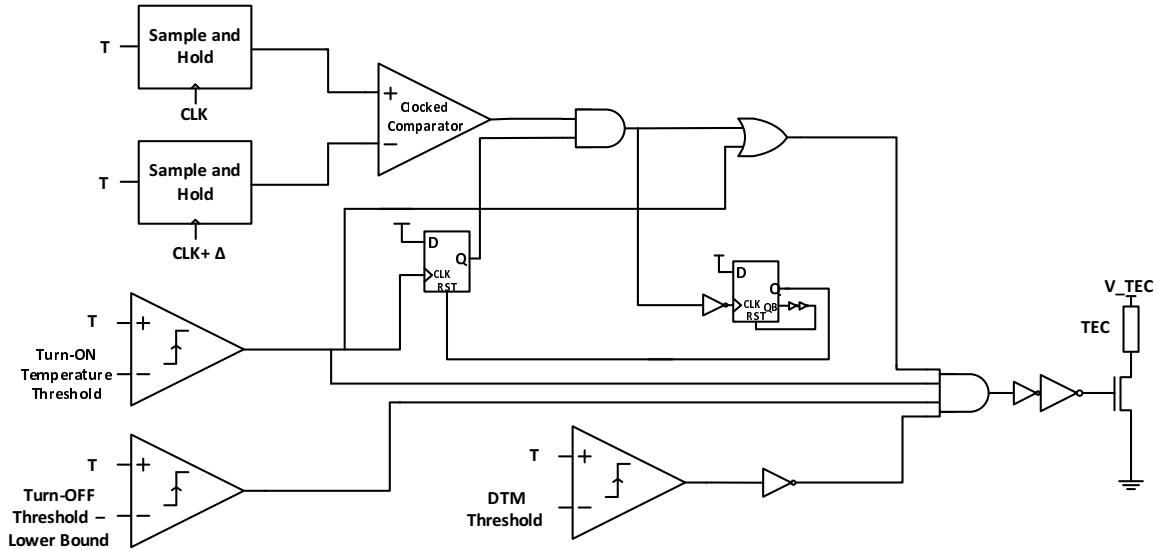
Figure 22: MCBC top level architecture.

4.3.2 Maximum Cooling Based Controller (MCBC)

Figure 22 shows a top level schematic of the MCBC. The controller is based on a sample and hold (S/H) architecture with two S/H circuits sampling the temperature sensor. The S/H circuits are designed using transmission gate switches that are sampling the input to a 100fF capacitor. At each time instant t, the output of the first S/H circuit is the positive input of a clocked comparator. A time shifted value $(t+\Delta)$ of the temperature sensor [with programmable delay (Δ)] is sampled by a second S/H circuit and the output is sent to the negative terminal of the comparator. Given this, the comparator should output a digital 1 if the more recent temperature T(t) is lower than the previously sampled T(t- Δ). This means that the TEC will remain on if the temperature is strictly decreasing. Once T(t) exceeds T(t-Δ) the comparator will switch and output a digital 0 turning the TEC off. Since we do not want the controller to turn on the TEC when the chip temperature is decreasing (during a low power event where cooling is not required), we would like the McBC to turn the TEC on only beyond a certain temperature. We have therefore added a threshold temperature detector that will only turn the TEC on once the temperature reaches a specified threshold using a positive edge detector circuit. Figure 23a shows a flow-diagram to explain the operating principle of the MCBC. This process will continue until the TEC cannot provide

38

any more cooling and the processor will have to be throttled by DTM techniques. We integrated the MCBC controller within our thermal package and simulated the response using the power pulse used in the ThBC simulations. As Figure 23b shows the temperature starts to increase initially until the threshold temperature is reached and the TEC turns on (simulations with HTC=625 W/m^{2}K). The TEC immediately begins cooling and the differential controller continues to keep the TEC turned on as the temperature is decreasing. Once the temperature reaches the minimum point the differential controller is turned off and resets the controller to monitor the temperature until it again reaches the threshold. The TEC again turns on and the differential control takes over until the temperature is strictly decreasing. As expected the cooling of successive dips becomes lower as the Peltier cooling of the TEC is reduced due to the Joule heating. Eventually the TEC cannot provide cooling when turned on at the threshold temperature, leaving the chip temperature to continue to increase until a given DTM technique is invoked to minimize the power dissipation. It can be observed that the higher levels of current provide more initial cooling but the larger Joule heating causes the cooling dips to have shorter duration. It can be observed that McBC reduces the number of on/off transitions of the TEC. If the temperature of the chip goes below the specified threshold the TEC will always turn off (irrespective of whether a local minimum has been reached) using a lower-bound comparator shown in Figure 22. This helps to prevent unnecessary energy dissipation in the TEC in the case that a high power pulse, that turned on the TEC and triggered the differential controller, disappears within a relatively short time-interval allowing the chip to cool naturally. The temperature response of the system for such a test case (i.e. a short duration high power event that triggers the hysteretic control) is presented in Figure 24. The test case considers a high power event until 45ms when the chip power goes to 0. As we can see the chip temperature is increasing and the McBC controller turns the TEC on. Shortly after the power goes to 0 and the temperature reduces due to natural cooling. The lower-bound threshold for the TEC is set at ≈67 °C for this simulation and it can be observed that the TEC turns off at
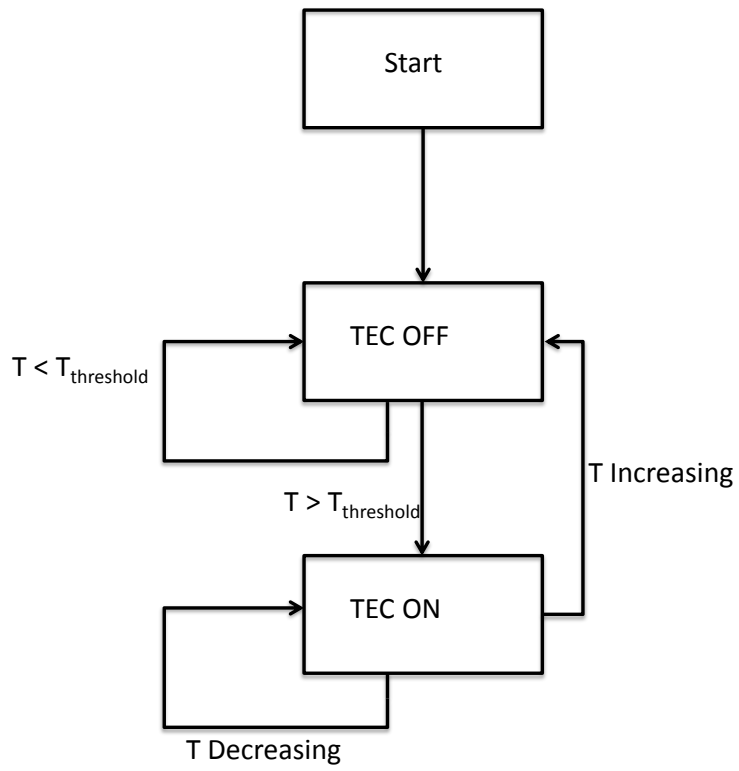
(a)
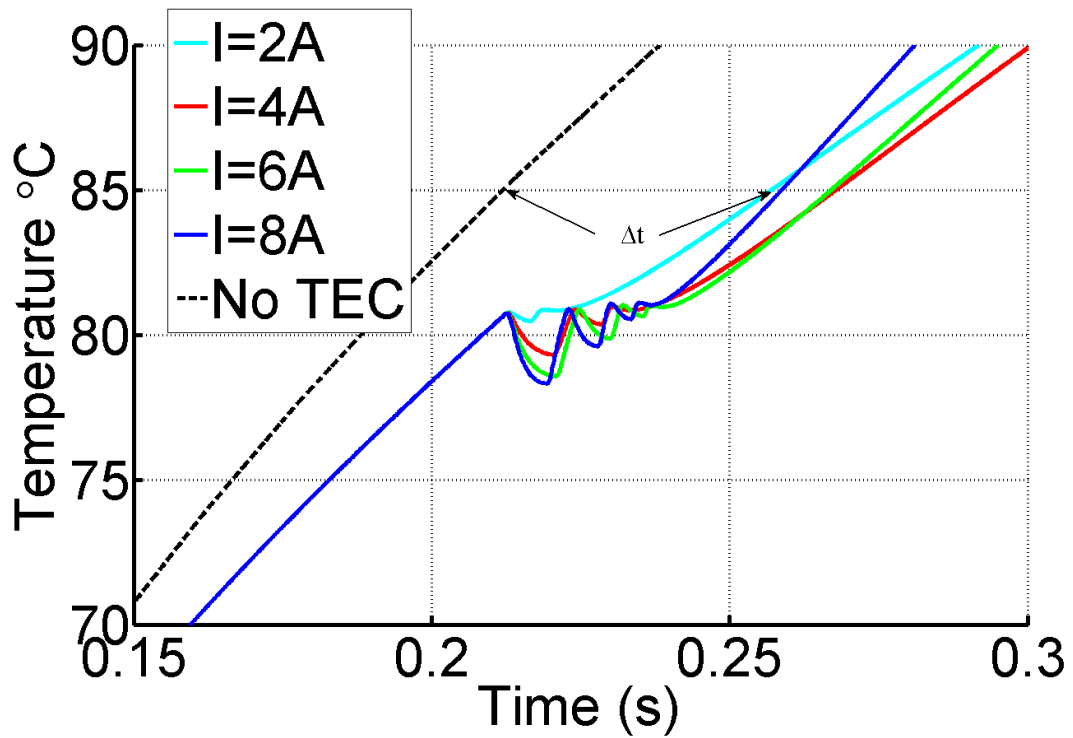
(b)

Figure 23: Maximum Cooling Based Controller: a) State Transition Diagram and b) simulation results.

Table 2: ThBC Characteristics.

ThBC Results					
TEC Current (A)	Δt (ms)	Number of Transitions	Average Temp ($^\circ$C)	TEC Energy (mJ)	Power FET Energy (mJ)
1	35.6	1	82.57	3.98	1.76
2	45.2	7	82.24	19.57	8.66
3	52.2	15	82.04	48.41	21.42
4	56.2	21	81.90	88.23	39.04
5	57.5	25	81.80	134.75	59.62
6	56.4	29	81.74	180.62	79.92
7	53.5	30	81.67	223.14	98.735
8	49.4	30	81.63	252.4	111.68
12	31	24	81.43	253.84	112.32

about 55ms. The temperature has a slight increase around this point as the Peltier current source has been turned off and the heat cannot flow through the TEC increasing the Silicon temperature. Beyond 70ms the temperature begins to decrease and eventually settles to a low steady state value. Once the TEC has exhausted its cooling ability and the temperature system invokes the DTM, the TEC should be turned off to save energy and allow the system to cool naturally under DTM. This can be accomplished using an additional comparator in both the ThBC and MCBC.

4.4 Controller Comparison

The overall characteristics of the two controllers are summarized in Tables 2 and 3 . The simulation results show that with the proper choice of TEC current, both the controllers can achieve ≈50-60ms of time extension which is quite significant at the microprocessor time scale. This time can be used to finish a workload without throttling and/or allocation to a different core in the chip. As explained earlier, increasing the TEC current initially increases the time extension due to more efficient transient cooling, but beyond a certain point, due to Joule Heating, the extension reduces. We further observe that the ThBC
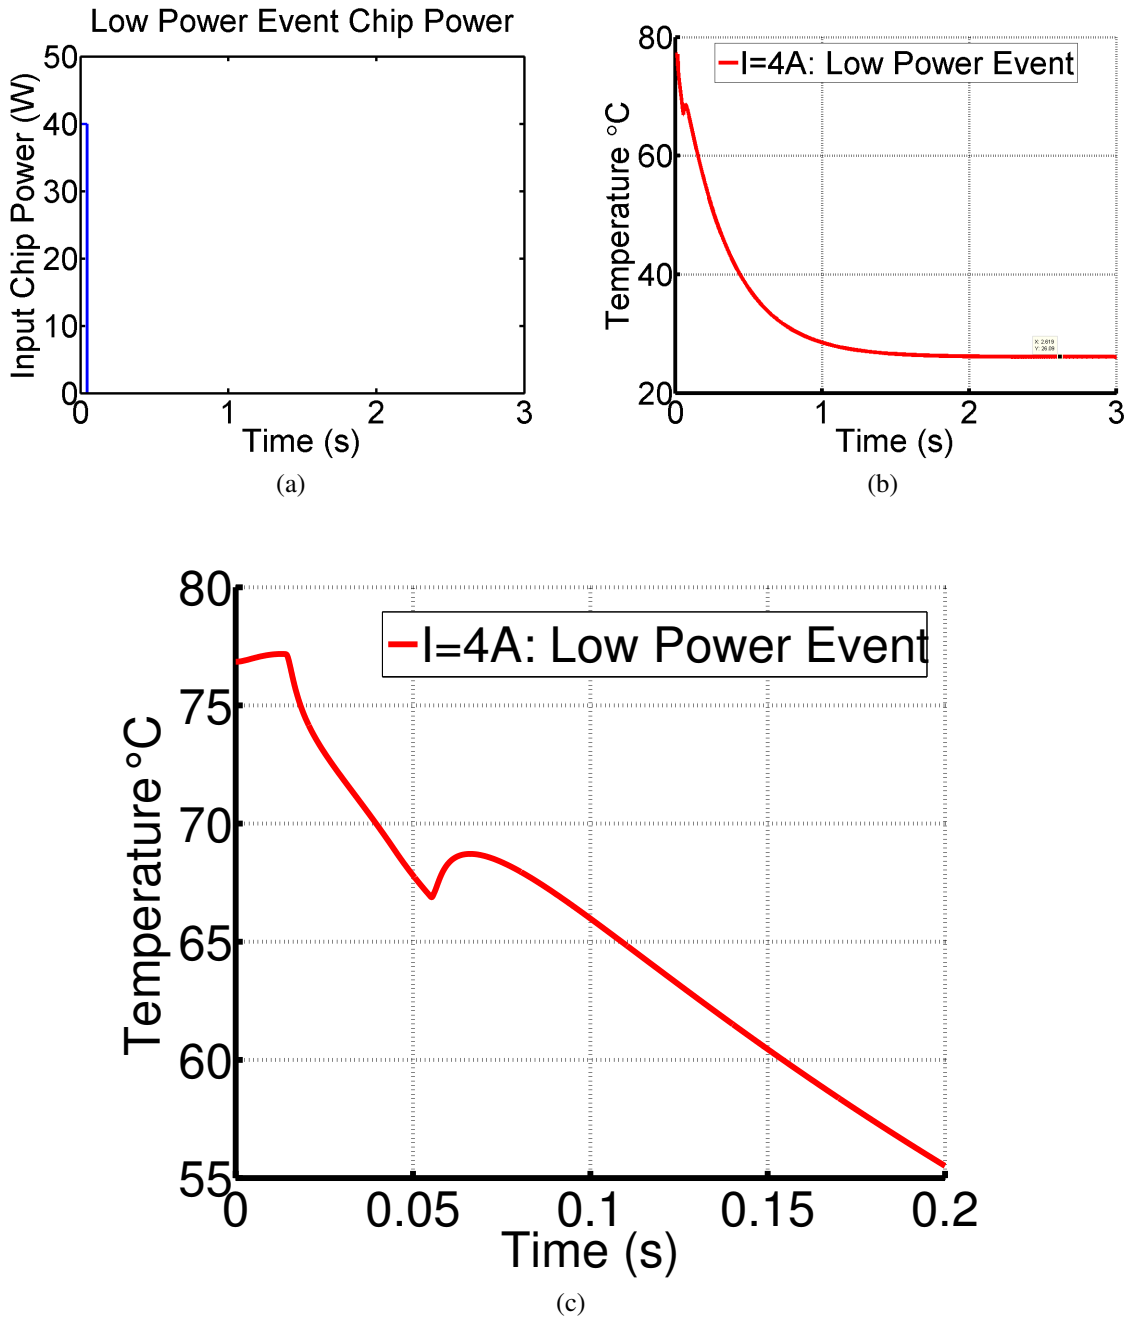
Figure 24: McBC with low power event: (a) the applied power pattern, (b) the observed temperature pattern, and (c) zoomed in view near 100ms to show the TEC operation. The power turns off at 45ms and the TEC shortly after at ≈5ms.

provides slightly longer extension times; about 1ms greater on average for a given TEC current. The trade-off is that the MCBC provides a lower number of transitions, 3-7 compared to the ThBC, which has about 7-30 transitions. The reduced number of transitions

Table 3: MCBC Characteristics.

MCBC Results					
TEC Current (A)	Δt (ms)	Number of Transitions	Average Temp ($^\circ$C)	TEC Energy (mJ)	Power FET Energy (mJ)
1	35.6	1	82.57	3.96	1.755
2	45.2	3	82.24	19.57	8.66
3	52.1	5	81.96	49.22	21.78
4	55.9	5	81.72	92.21	40.08
5	56.7	5	81.51	145.21	64.25
6	55.1	7	81.41	199.33	88.2
7	51.6	7	81.32	252.49	111.72
8	46.9	7	81.30	296.51	131.2
12	25.9	5	81.24	361.24	159.84

can improve the reliability of the TEC, as fewer transitions could extend the TEC lifetime [42]. The MCBC also results in a lower average temperature during the power pulse, about 0.3 $^\circ$C on average when compared at same current levels. This reduction in temperature can reduce the overall leakage power on the chip.

Next we characterize the energy overhead of each controller. To accurately estimate the overhead, a piece-wise linear curve of the temperature profile of the chip as shown in figures 21 and 23b are applied to the appropriate controller. The energy overheads originate from three main sources: (a) the operating power of the controller including the switching energy dissipation while driving the gate capacitance of the power FET; (b) the power loss in the TEC (I^2R_{TEC}); and (c) the loss due to the finite on resistance of the power FET during on period (I^2R_{ON}).

The varying levels of switching activity within the controller circuits do not contribute much to the energy dissipated in the controller as most of the energy is in the analog blocks biasing currents. Since the MCBC has more analog control blocks (mainly the comparators), the biasing (static) energy is higher for the MCBC. A higher number of transitions in the TEC tend to add additional switching energy due to the gate capacitance of the power

FET and the output capacitance at the intermediate node between TEC and FET, but the switching energy has negligible contributions to the overall controller energy in both the ThBC and MCBC. However, as shown in Table 4 , the controllers energy is almost insignificant when compared to the TEC energy. The power dissipation of the TEC increases significantly with increasing levels of TEC current. This suggests the choice of the TEC current is a trade-off between the time extension and energy overhead associated with cooling. The analysis of the energy overhead as a function of the control principle and TEC current is an interesting problem. The maximum energy overhead of the TEC occurs when the TEC is always on and continuously consumes current. The "ON" power of the TEC is defined as:

$$P_{TEC} = I_{TEC}^2 (R_{TEC} + R_{FET}) \qquad (2)$$

The maximum energy overhead can be computed as:

$$OV_{MAX} = \frac{P_{TEC}}{P_{CHIP}} \qquad (3)$$

and is summarized in Table 5. The equivalent resistance of the TEC is approximately 113mΩ and includes the resistance of the contacts as well as the super-lattice material. The power FET was designed for approximately 50mΩ ON resistance in a 130nm CMOS process and had an area of 600μmX600μm in layout, about 5% of the total TEC area. For example, the peak chip power dissipation in this analysis was 42.4W, for 4A current the maximum energy (or power) overhead of the TEC will therefore be 6.2%. A more accurate

Table 4: Controller Energy Analysis.

Controller Energy Analysis		
TEC	ThBC	MCBC
Current (A)	Energy (μJ)	Energy (μJ)
2	0.550	1.978
5	0.526	1.898
8	0.528	1.980
12	0.546	2.102

analysis of the energy overhead needs to consider the fact that the TEC may not be always "ON" and it goes through "ON" and "OFF" cycles depending on the control principle. Assume a power pulse of duration T which results in a total chip energy

$$E_{CHIP} = P_{CHIP}T \qquad (4)$$

The energy consumed by the TEC for the same time duration of T is given by

$$E_{TEC} = E_{\Delta t} + I^2_{TEC}(R_{TEC} + R_{FET})(T - \Delta t) \qquad (5)$$

where, $E_{\Delta t}$ is the energy dissipated during the extension time of Δt. The second component of the preceding equation defines the energy dissipation of the "ON" TEC beyond the extension period. The energy overhead for a duration T is defined as

$$OV(T) = \frac{E_{TEC}(T)}{E_{CHIP}(T)} \qquad (6)$$

Table 5 summarizes the analysis of the overhead considering a 100ms power pulse. The TEC current and the control principles impact both Δt and $E_{\Delta t}$ (as illustrated in Tables 2 and 3) and hence, modulate the energy overhead. First, the extension time Δt depends on the current and control principles; and second the energy dissipation during the extension time $E_{\Delta t}$ depends on how long the TEC (and FET) remains "ON" in this period which in turn depends on the TEC current and the control principle. As explained earlier, Δt initially increases with TEC current but start to reduce beyond a certain point. However, due to the quadratic relation between the TEC/FET energy and the current through the TEC, we observe that the overall energy overhead is always increasing with increasing TEC current (Table 5). During the extension period, the MCBC keeps the TEC on for longer a period of time and hence, results in higher energy dissipation in the TEC and FET compared to the ThBC (Table 5). Further, the MCBC also has a marginally lower extension time compared to ThBC. Therefore, we observe that for a given TEC current, the MCBC consumes marginally more energy than the ThBC even with less number of transitions. Figure 25 graphically summarizes the trade-off associated with the choice of the

Table 5: Controller Energy Overhead.

TEC Energy Overhead			
TEC Current (A)	Maximum Overhead (OV_{MAX})	Overhead with 100ms Power Pulse of 42.4W (OV_{MAX})	
		ThBC	MCBC
1	0.4%	0.4%	0.4%
2	1.5%	1.5%	1.5%
3	3.5%	3.3%	3.3%
4	6.2%	5.7%	5.8%
5	9.6%	8.7%	9.1%
6	13.8%	12.2%	13.0%
7	18.8%	16.4%	17.7%
8	24.6%	21.0%	23.2%
12	55.4%	46.8%	53.3%

TEC current for the time-extension and the energy overhead for both the control principles. As observed in the figure, both the control principles have a similar trade-off between the TEC current and time-extension. The choice between the two control principles depends on the marginally higher extension time and lower overhead for ThBC versus the reduced number of transition and marginally lower average temperature in MCBC.

4.5 Effects of Parasitics on Controller Design

As observed each controller generates a digital control signal that has to drive a power FET which modulates the current within the TEC. The control signal turns the power FET switch ON and OFF and allows the current to flow into the device. The current level is modulated from an external supply voltage which can also be digitally controlled from inside the chip. Since the controller has to drive an external TEC it is very important to consider package and board parasitics when evaluating the performance of the controller. The parasitics can severely impact the cooling performance of the TEC, as they can reduce the current through the TEC as well as increase the response time. It is critical for a designer to account for these parasitics and design a proper current drive solution to mitigate the performance
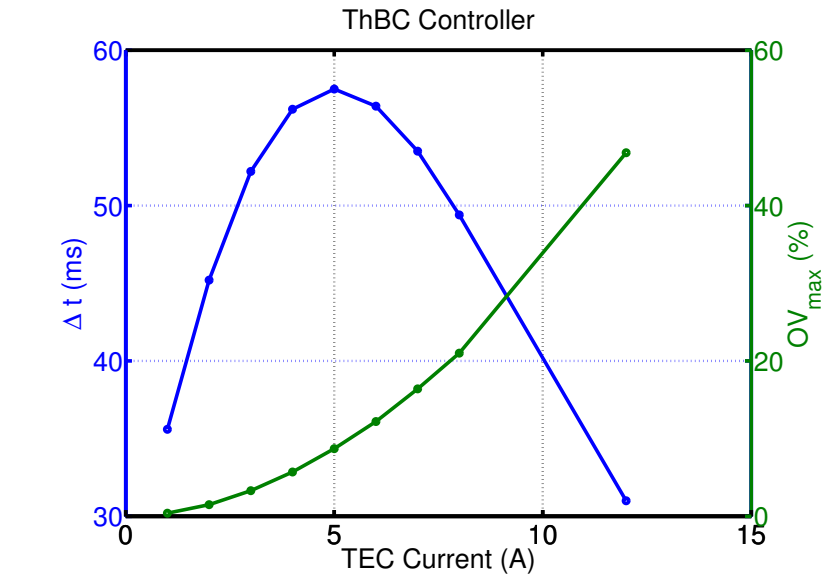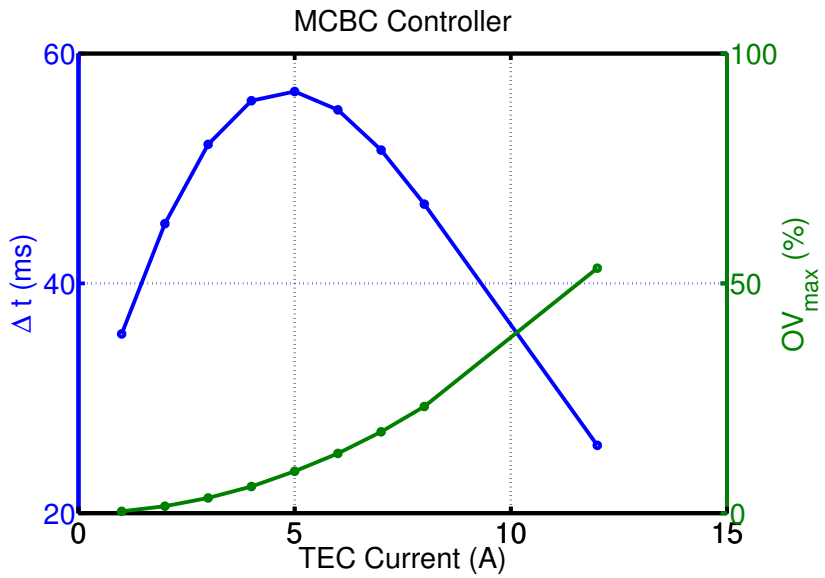
Figure 25: The trade-off analysis for the proposed control methods: (a) ThBC and (b) MCBC.

degradation. In order to assess the effect of the parasitics on the cooling solution we focus on 2 specific cases: 1) TEC control with power FET on chip, and 2) off-chip power device. We discuss the trade-offs between each solution.
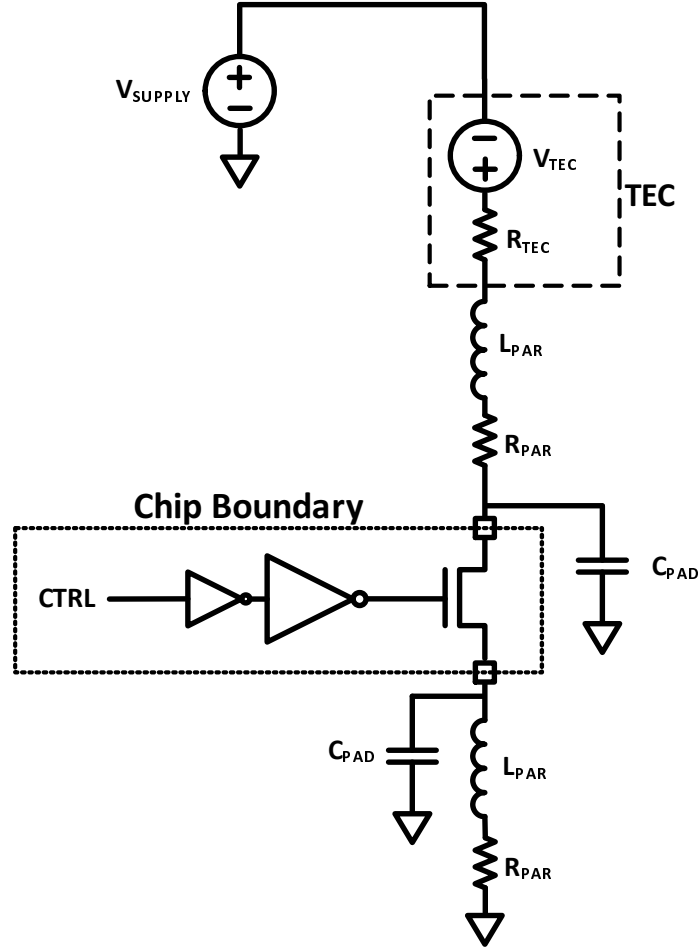
Figure 26: Schematic of the on-chip controller and power FET solution. The bondwire and board level parasitics are included to study the performance degradation.

4.5.1 On-Chip Power FET

Figure 26 shows the on-chip controller framework with a generic control signal driving an on-chip power FET. The CTRL signal is generated from a control principle that senses the temperature of the silicon and makes a decision whether to turn the TEC ON or OFF such as the previously discussed ThBC and MCBC. We drive the gate of the power FET and turn the MOSFET on and off. The source and drain of the transistor are connected to

chip pads. These pads have a capacitance C_{PAD} that is on the order of 10 pF for modern packages. The pads are then connected out from the chip using bondwires to the board level. The bondwires have inductance in the order of 1-10 nH and resistance on the order of 200mΩ. We model these as L_{PAR} and R_{PAR}. Note that the resistance includes the effects of the bondwire and the board trace resistance. An external voltage supply supplies the TEC current with the power FET turning the TEC ON and OFF. We can clearly see that the current that flows through the TEC will be determined by the total series resistance seen

$$I_{TEC} = \frac{V_{SUPPLY}}{R_{TEC} + 2R_{PAR} + R_{FET}} \tag{7}$$

and each bondwire will severely limit the total current. In addition the FET has a maximum voltage that can be applied to its gate and drain before the gate oxide breaks down. In the IBM 130nm technology the regular *nfet* has a maximum gate to source breakdown voltage of 1.4V. This means that for a TEC device with $\Delta T = 10°C$, (typical operating condition of our device) we will have a voltage of around 300mV across the TEC. This means that when the FET on chip has the gate grounded, the drain will have the sum of the supply voltage and the TEC voltage across it. This means that V_{SUPPLY} can only go up to 900mV. We have left some margin due to noise that can occur during fast switching of the transistor. In addition, we do not want to put a very large transistor on chip to avoid area overheads. We have limited the size of our transistor to 0.1mm^2 which accounts for about a 1% area overhead for a 3mmX3mm TEC.

Figure 27 shows the simulation results current through the TEC vor varying V_{SUPPLY}. We have used transistors of layout area 10000μm^2, 40000μm^2, and 80000μm^2, which is roughly an overhead of 0.1%, 0.4% and 0.8% of the TEC area. We use values for the parasitics as follows: C_{PAD}=10pF, R_{PAR}=250mΩ, L_{PAR}=5nH. These values are consistent with the packages used by MOSIS with a typical 1mil bondwire. The results show that increasing the size of the FET has marginal improvement on the maximum current due to the decreased ON resistance. However the current is mostly limited by the maximum value of the supply voltage and parasitics resistance. We conclude that with the regular *nfet*

device, the maximum current is 1.83A. Based on the results of the ThBC and MCBC, this can severely limit the cooling that the TEC can provide, as those controllers had optimum current between 4 and 6A. Our transient simulation with the largest transisor shows that the TEC current turns on 203ns after the control signal is asserted, which includes the capacitance of the MOS gate and pad. This is fast enough to respond to the proposed ThBC and MCBC controllers. Additionally there was no overshoot due the parasitics inductance with this rise time.
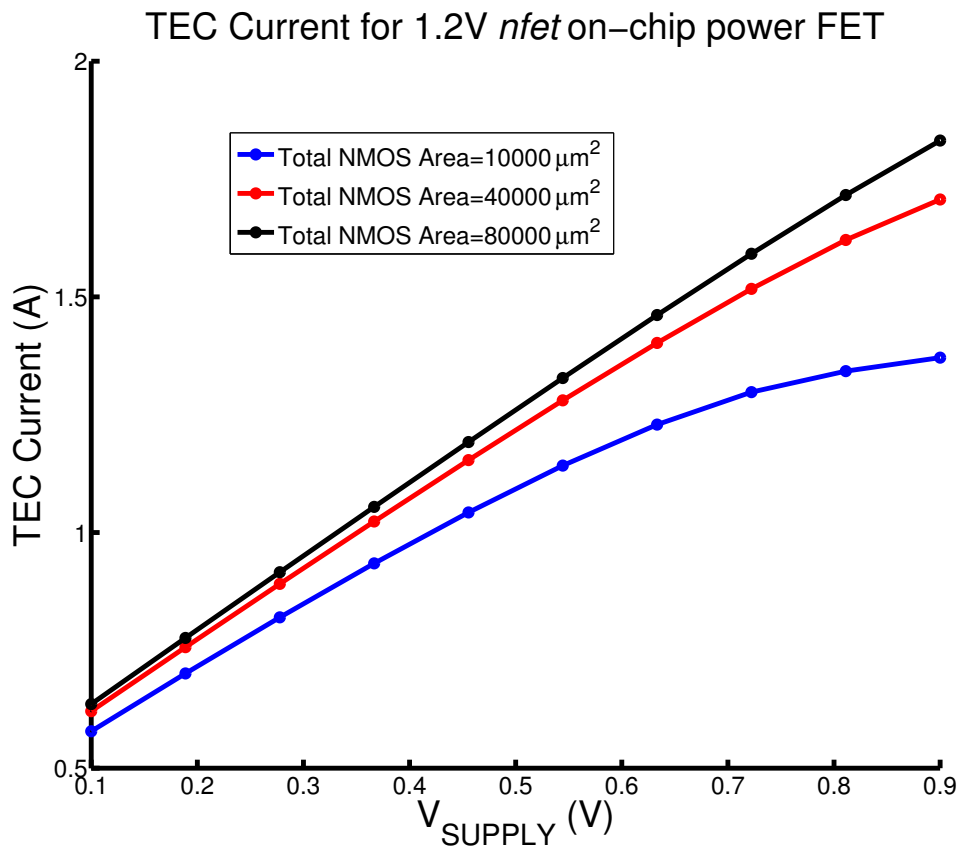
Figure 27: Simulation results showing the effect of the parasitics on the maximum current using a nominal nfet device. The extra resistance can severely limit the maximum current that can be sourced and even with the largest transistor, we can only source 1.83A.

In order to increase the maximum current we need to use the higher voltage FETs typically provided in modern processes. In the IBM 130nm process we have the *nfet33* which is a high voltage device that operates with supply voltages up to 3.3V. This means

that the device can sustain higher gate to source voltages before oxide breakdown. This means that the maximum that V_{SUPPLY} can reach without causing oxide breakdown is 2.9V. The drawback is that the transistor has larger ON resistance for the same size as compared to the normal *nfet* device. Figure 28 shows the simulation results using the on-chip high voltage device. We use the same parasitics as in the *nfet* simulation. The maximum current that can be applied to the TEC is now 4.58A. This is closer to the optimum level for the ThBC and MCBC, but is still limited. Our transient simulation shows that the current turns on 55ns after the control signal is asserted, which includes the capacitance of the MOS gate and pad. Again this rise time is fast enough, as both our ThBC and MCBC only require sub 1ms turn-on times. Like the 1.2V device, the parasitic inductance did not cause any overshoot or ringing.
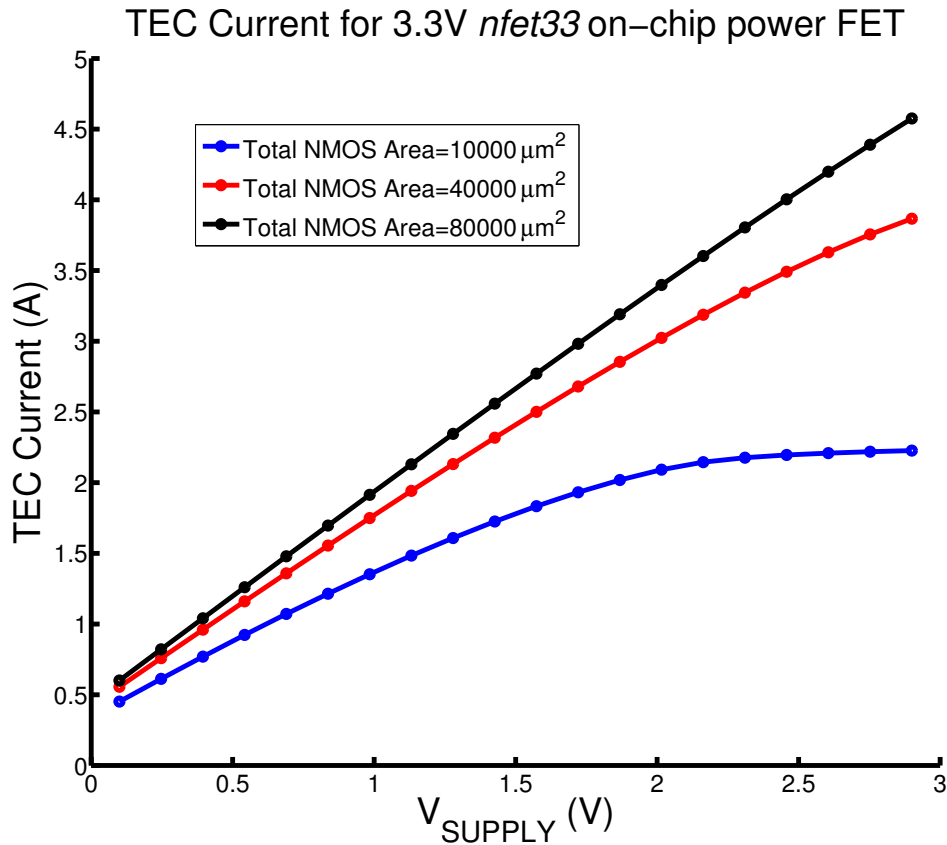
Figure 28: Simulation results showing the high voltage nfet used as the power FET. The TEC current is limited to 4.58A using the largest transistor.

Lastly the bondwires have a maximum current that they can handle before they melt. This is called the fusing current. For a 1mil diameter gold bondwire the fusing current is around 1.8A. This means that if the on-chip power FET has a single pad connection to the TEC, the current will be limited by this fusing current as we do not want to melt the wire. We can get around the fusing current limit by using multiple pads for the source and drain connections to provide multiple bondwires and reduce the parasitic resistance, but since many of today's chips are pad limited, a designer might not want to sacrifice multiple pads for the TEC. We therefore need to further consider an off-chip solution.
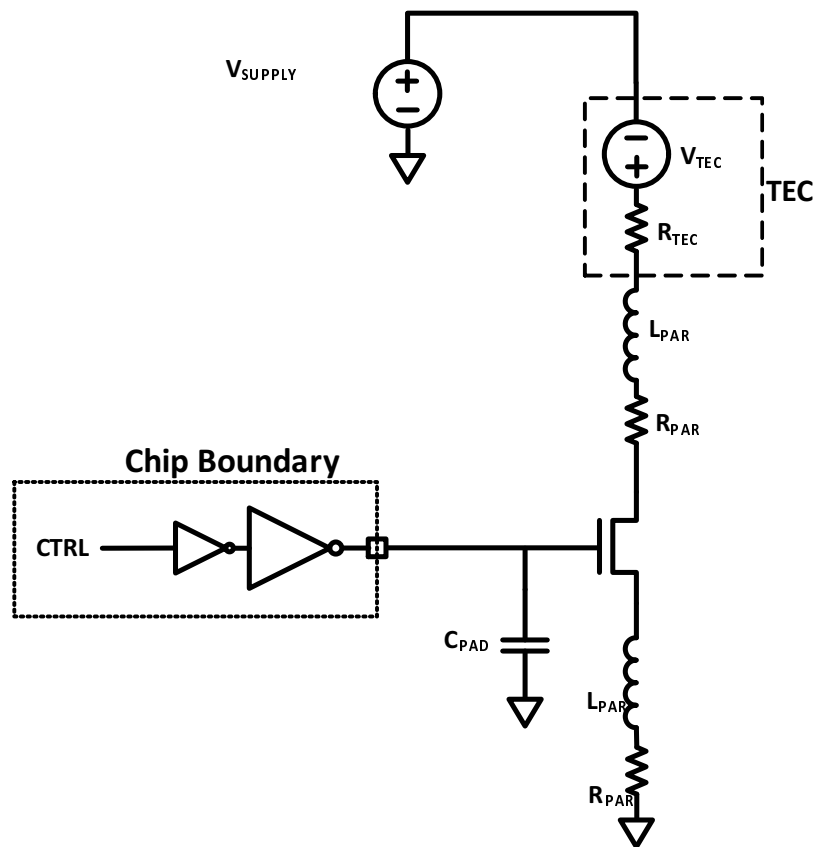
Figure 29: Schematic of the off-chip controller and power FET solution. A benefit of this solution is that only a single pad needs to be driven off-chip.

4.5.2 Off-Chip Power FET

Figure 29 shows the off-chip controller framework. The only difference from the previous implementation is that the chip drives a digital signal out that gets level converted off-chip

and drives a power FET device with its own package. We can see that the control signal now only needs to drive the gate of the off-chip transistor and now requires only 1 pad and its associated total capacitance which included the pad capacitance and the gate capacitance of the MOSFET. We can simply drive the pad with the 1.2V logic without the need of a level shifter, as devices like Alpha and Omega Semiconductor's AON2400 NMOS [43] can provide 24mΩ ON resistance at 1.2V of gate drive. The device can deliver a current up to 8A with 8V at the drain of the transistor. Transient simulation results show that the simulated transient control voltage at the gate of the off-chip transistor using the same driver that was used to drive the on-chip power FET, has a rise time of around 7.52μs considering the sum of the pad and transistor gate capacitance (1.7nF) as well as gate resistance (4Ω). This means that we can provide turn-on times of less than 1ms which is fast enough for the thermal transient extention times described by the ThBC and MCBC. The only drawback of the off-chip solution is that it requires extra board space to fit the package of the power FET. The device that we use has a footprint of 2mmX2mm which can add significant overheads to the design of the board. However considering the effects of the limited current levels that the on-chip solution suffers, the off-chip solution is recommended.

4.5.3 Recommendations

The analysis above has confirmed that package constraints and parasitics can cause significant reduction in the maximum TEC current that can be sourced using an on-chip solution while requiring a larger number of pads to access the transistors and provide enough current drive, so we conclude that it is better to drive an off-chip power-device and only keep the controller logic on-chip with 1 pad connection. This ensures maximum cooling performance by sourcing the optimally high 4-6A of TEC current with the only drawback being increased board level area. For a high performance microprocessor or server system this is a viable trade-off as the board design in typically not space limited.

4.6　Effects of Cooling Solution on TEC Performance

It is intriguing to consider the effect of the cooling solution on the TEC performance. Heat sink solutions for microprocessors have different sizes and fluid flow rates, which account for changes in the effective amount of heat that can be removed from the chip (external cooling). This is modeled with an effective Heat Transfer Coefficient (HTC), which varies from 500 W/m^{2}K to 20,000 W/m^{2}K [35]. The cooling solution can have a significant effect on the total TDP of the package and as such can restrict the maximum allowable power that the chip can sustain. Since the power profile will have an impact on the TEC performance, we study the effect of different ranges for the HTC on the controller performance. Figure 30 summarizes the effect of varying heat transfer coefficient (HTC) from the chip on the effectiveness of the TEC assisted cooling. As observed, the initial steady state temperature is higher for the lower HTC case. When the power pulse is applied, the temperature increases at a faster rate with a lower HTC and the threshold temperature is reached earlier. Moreover, with a higher HTC, the controller is able to sustain the threshold temperature for a longer period of time and hence provide a longer time extension for both ThBC and MCBC. Further, in the case of MCBC, once the cooling begins we observe deeper dips in the temperature profile with higher HTC. Therefore, we expect a lower average temperature with MCBC at higher HTC. Figure 30 shows that a higher HTC results in a larger time extension for both the ThBC and MCBC. This proves that the TEC is effective across all external cooling solutions, but is benefited from having a better external cooling solution.

4.7　Simulation of controllers considering architecture and workload

We verify the TEC assisted cooling considering processor workload driven power estimates. The power trace was generated via architectural simulations to reflect the characteristics of executed workloads and microarchitecture. The SPEC2006 benchmark suite was executed with a detailed cycle-level x86 timing simulator, Zesto [44]. Zesto was configured to model a 4-issue out-of-order pipeline and generate access counts of architectural
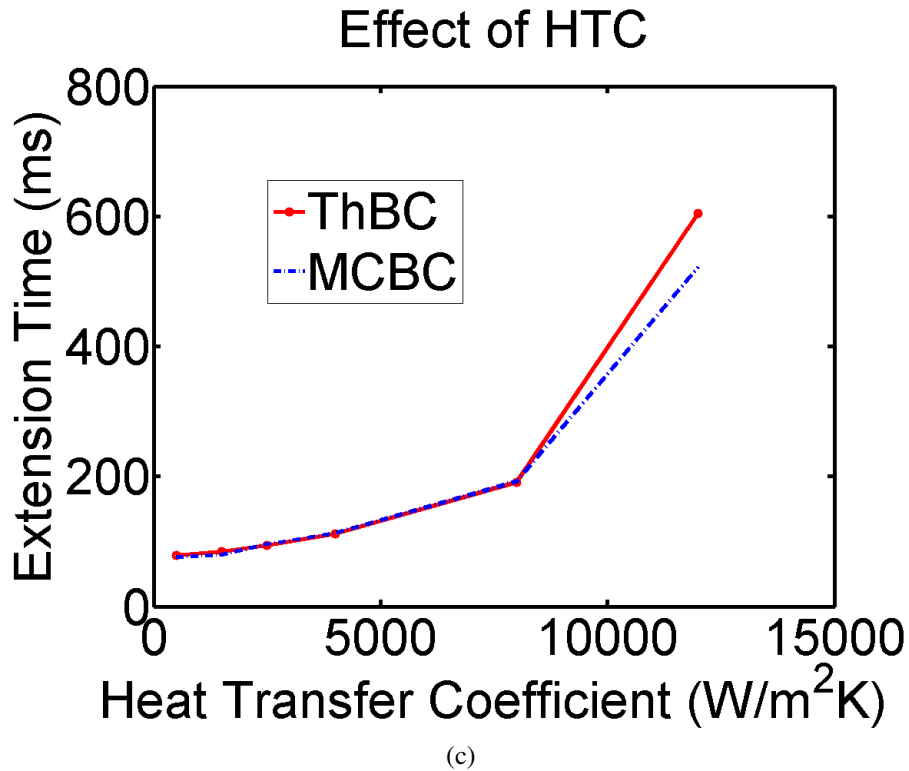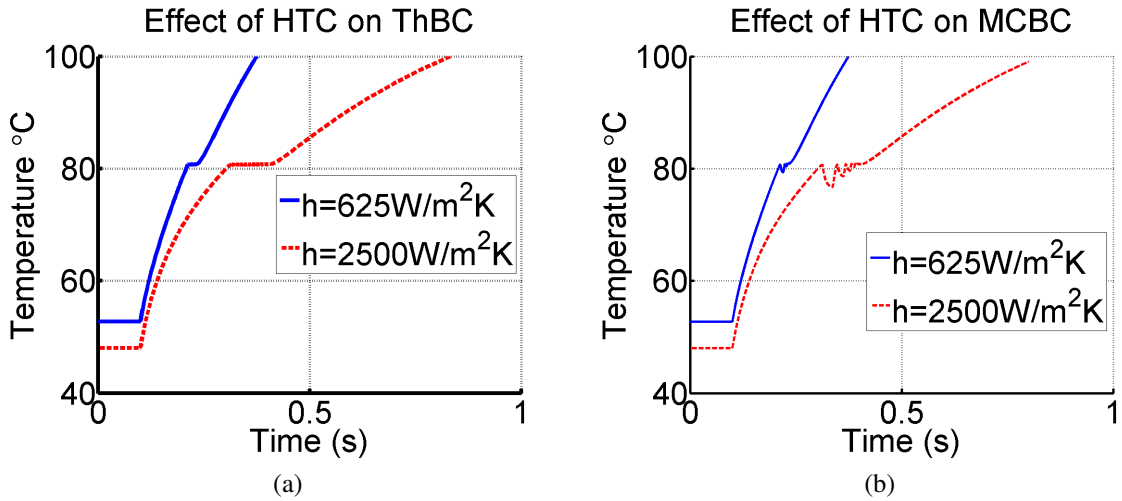
Figure 30: The effect of heat transfer coefficient (HTC) on the controller operation: (a) effect of HTC on ThBC, (b) effect of HTC on MCBC, (c) effect of HTC on ThBC and MCBC time extension.

components. The access counts indicate how the workload exercises different architectural components. In conjunction with Zesto, McPAT [45] was used to estimate the energy dissipation of architectural components. Dynamic energy is calculated by multiplying access
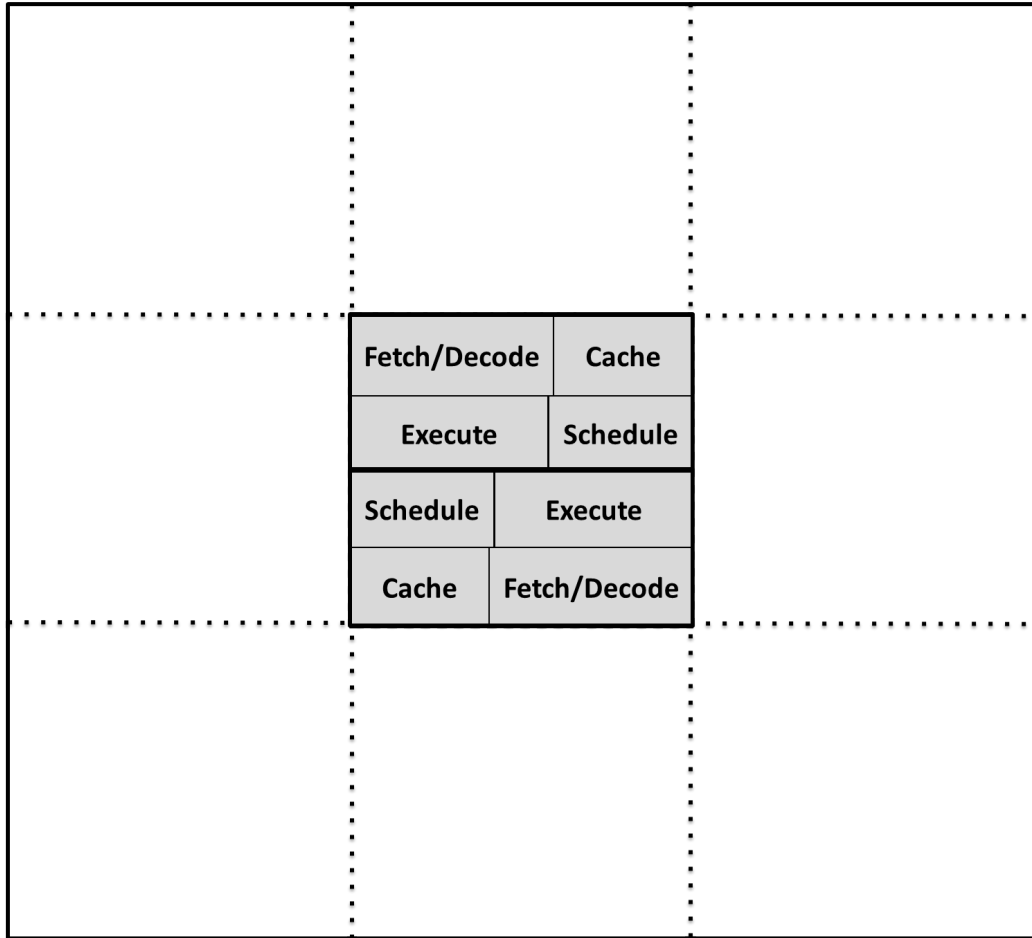
Figure 31: Chip Floorplan.

counts with estimated per-access energy, and the total power is the sum of dynamic and leakage powers. A related work [46] used Zesto and McPAT with SPEC2006 benchmarks to validate the accuracy against Intel cores. The die floor plan, was made based on the area estimation provided by McPAT and has two cores underneath a 3x3mm TEC centered on a 9x9mm chip. Each core is partitioned into blocks by pipeline stages and cache. This is then converted to a grid using the exact sizes of each block and appropriate floorplan, and inserted under the TEC area. We have added 10W of background power elsewhere in the chip and ran the same simulation. The benchmark we used was *sjeng* (a relatively high power workload) and it was simulated on a 1.2V-OOO core implemented in a predictive 16nm technology.

Table 6: Architectural Level Simulation Paremeters.

Configuration	Description
Instruction set architecture (ISA)	Intel x86 IA32
Core Pipeline	Out -of-order execution
Pipeline Width	4-wide pipeline (6 peak issue width)
ROB Size	128 entries
Instruction/L1 Data Cache	32KB, 4-way associative, 32B line
L2 Cache	256KB, 8-way associative, 64B line

The surface plot in Figure 32a shows the 3-dimensional view of the power density (W/cm^2) at a time instant (t=100ms) across all elements (each element has a size of 1mmX1mm) in the chip. It can be observed that the workload creates a hot spot in the cores. Figure 32b shows the same power profile across a top level view of all the elements. Figure 32c shows the input transient power variation across the center and edge of the TEC, as well as the chip edges where background power was applied. Figure 32d shows that the TEC assisted cooling allow us to maintain the hotpot temperature (the maximum temperature of the chip) below the threshold for a longer time period. We can see that no matter what controller we use, we can manage the transient temperature below a threshold for a significant amount of time and extend the benchmark workload beyond the time of thermal limit. This is additional performance for the system.

The above analysis validates that the proposed control principles works with a realistic processor workload and motivates future work on the co-design of the TEC control principle and run-time power management techniques. Coupling the architectural level development with the TEC can lead to significant performance gains and needs to be co-developed.
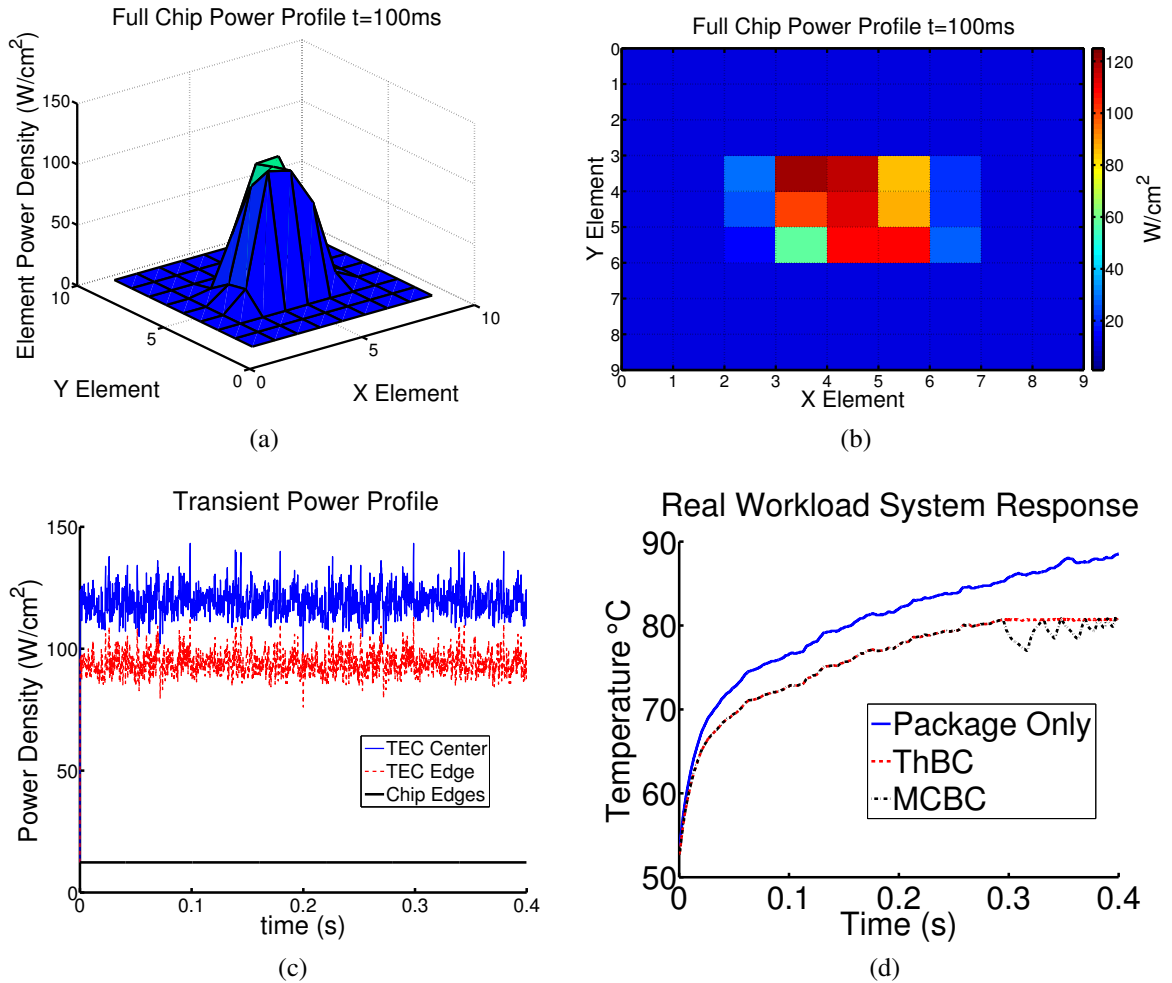
Figure 32: Analysis of TEC assisted cooling with a processor workload (a) 3D view of the full-chip power density pattern at t=100ms, (b) Top level view of the power profile at t=100ms, (c) Time domain power density variation across sections of the TEC and background elements, and (d) transient temperature variation for the package, and the ThBC and MCBC at the TEC center.

4.8 Summary

In this chapter we have studied the prospect of on-demand cooling with super-lattice thin-film TECs integrated in a chip and package. We have demonstrated on-chip controllers for temperature dependent dynamic activation/deactivation of TECs to provide on-demand cooling. The electro-thermal analysis was performed integrating the transistor level models of the control circuits, designed in 130nm CMOS, with full-system thermal compact model of chip, package, and embedded TEC. The co-analysis shows that TEC assisted transient

cooling allows a processor to sustain a high power pulse for a longer period of time without violation of the thermal threshold. This can significantly reduce the thermal events in processors and help improve thermally limited system performance, which will become critical in future generation systems. The possible performance gains of integrating TECs within a microprocessor package suggest the need for future wok on this topic on designing more efficient controllers as well as on innovative approaches to exploit TEC assisted cooling through micro-architecture driven DTM approaches. The potential advantages of integration of TECs can also inspire investigations in chip-package co-design.

CHAPTER 5

ENERGY-EFFICIENT AUTONOMOUS ENERGY MANAGEMENT SYSTEM

5.1   Introduction

A major concern for TEC based on-demand cooling is the need for additional energy. The processors experience dynamic variations in power dissipation during operation. The TEC assisted cooling is necessary only during thermally critical high power modes; and the TECs are normally turned off during nominal power modes. The finite heat flux generated during the nominal power modes is wasted as the heat energy is dumped in the environment. It is intriguing to note that this heat flux flows through the TEM and can be harvested to generate electrical energy by operating the thermoelectric material (TEM) in the Seebeck mode as a thermoelectric generator. Figure 33 shows this concept. The thermal events occur with very low frequency, and therefore even though the TEC is providing natural cooling due to higher thermal conductivity of the copper and thin-film super-lattice material, most of the time the TEM is not being utilized for active cooling. Being able to take advantage of the chip's idle power and the heat that is generated, we can increase the overall energy efficiency of a cooling system on chip. The switching of the TEM to the TEG mode allows a part of the otherwise wasted heat energy to be reclaimed and stored for later use. The stored energy can be used to provide cooling during the intermittent high power modes or even power sections of the chip.

The TEG operates on the basis of the Seebeck effect where a temperature difference across a device creates an open circuit voltage as shown in equation 8. This is due to the natural movement of the electrons and holes in the p and n poles of the device due to the temperature gradient, creating the open circuit voltage.

$$V_{OC} = S(T_H - T_C) \tag{8}$$

Since the difference between T_H and T_C can be small during nominal power modes, the
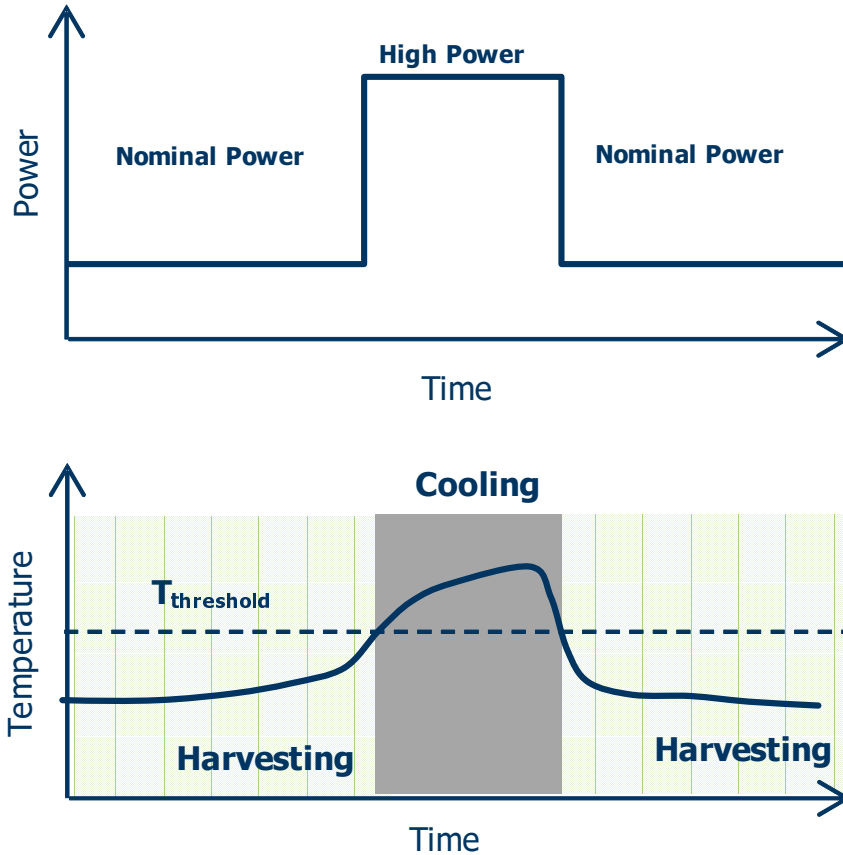
Figure 33: The overview of the proposed system. In nominal or idle power modes, the system operates in the harvesting mode and acts as a TEG storing energy. In high power modes when active cooling is required, the system moves into the cooling mode and helps mitigate thermal issues.

voltage needs to be boosted to a usable level. This can be done with a boost converter. The thermoelectric material also has an electrical equivalent resistance due to the internal Joule heating as well as the copper contacts. This can significantly reduce the effective power that can be harvested from the TEG. The TEG can therefore be modeled as a simple thevenin circuit equivalent. Figure 34 shows the model.

The goal of this chapter is two-fold: to design and experimentally verify a system that is able to use a single thermoelectric device for both harvesting and cooling, and also verify the control methods and TEC cooling developed in the last chapter. The system is fully on-chip and is designed to connect to an external TEM. We present a fully integrated on-chip system for energy-efficient on-demand active cooling using dynamic mode switching of a
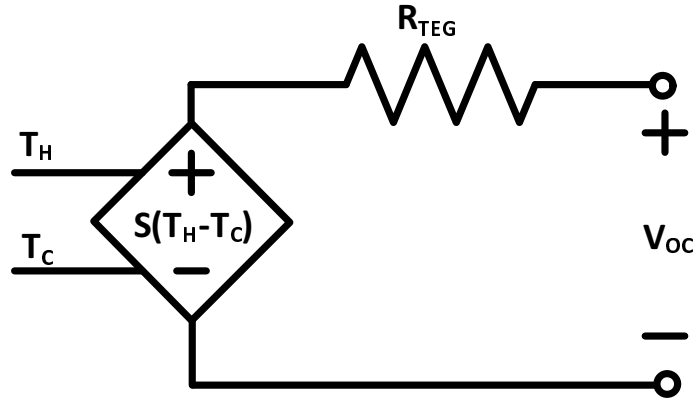
Figure 34: The equivalent electrical circuit model for the TEG.

single TEM. The proposed system includes a high-efficiency boost regulator encompassing low power design techniques to harvest heat energy and store it in an off-chip capacitor. This capacitors stored energy is then used to power a constant current source when cooling is required. The current source can deliver the desired current to the TEC over a wide range of supply voltage. The current source can be programmed digitally to change the cooling current and hence, degree of the achievable cooling. Once the capacitor energy is exhausted, the system automatically switches to powering the cooling from the chips supply voltage. A comparator and switch matrix control the mode switching of the TEM. A test-chip of the proposed system is designed in 130nm CMOS. The measurement of the test-chip demonstrates the system functionality. The test-chip is tested with a commercial TEM to demonstrate the dynamic mode switching and simultaneous cooling and energy harvesting.

5.2 System Level Overview

Figure 35 shows the proposed top level system. The boost regulator boosts the voltage generated by the TEM and dumps charge in an output capacitor. This is called the harvesting mode. Harvesting is active at all times during which the silicon temperature is below a specified user reference, and the TEM operates in the Thermoelectric Generation (TEG)

mode. The mode switching is done using a switch matrix, which is responsible for connecting the TEM terminals in the appropriate fashion. When the TEMs negative terminal is grounded and the positive terminal is connected to the inductor, the TEM operates in the harvesting mode by the Seebeck effect. In order to harvest energy, we must connect the positive TEM terminal to the inductor and the negative terminal to ground, so we turn transistors M1 and M4 on, while keeping M2 and M3 off. In the cooling mode, current is forced into the negative terminal of the TEM, making it operate in the Peltier mode and dumping heat from the hot to the cold side of the TEM. Our system does this by turning switches M2 and M3 on and turning M1 and M4 off. The switching logic is implemented in the TEC mode controller and is discussed in the following section. The control signals $\Phi 1$ and $\Phi 2$ control the switch matrix transistors. The system only has 2 modes of operation and is always in either the cooling or harvesting mode. During low power events, when the temperature of the silicon is not high, the system will stay in the harvesting mode and dump energy to the output capacitor. This energy is stored in the output capacitor and can be used to power the TEC controller. Once the capacitor gets charged and the voltage reaches the regulation point of the boost regulator, no more energy can be harvested
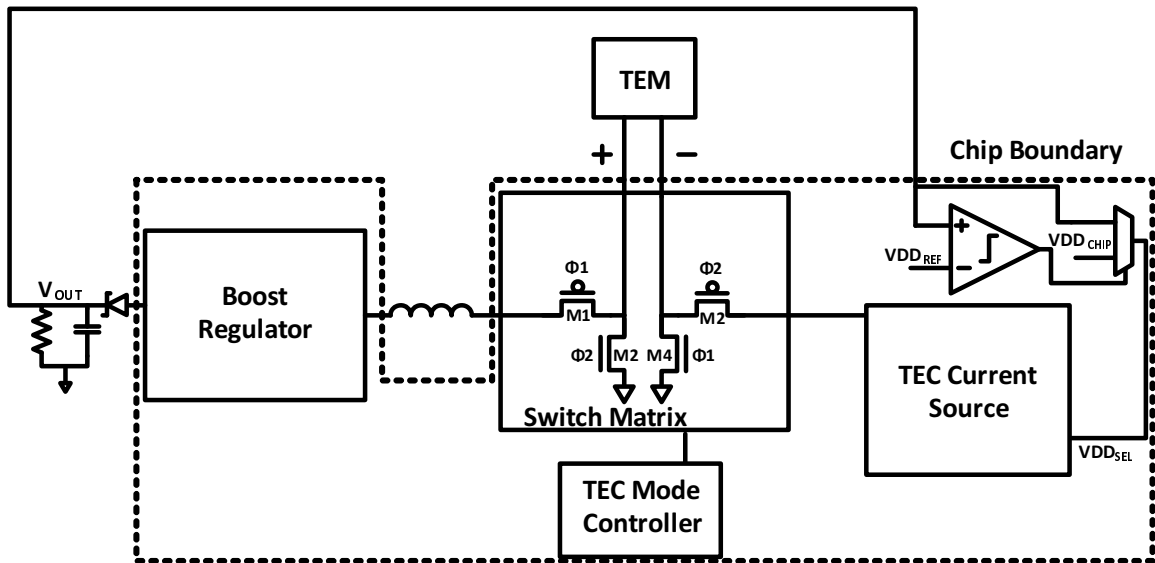
Figure 35: Schematic level overview of the proposed integrated TEM control

by the TEM. Although this is an issue for long-term operation, a battery charger could be implemented in mobile systems to recharge a battery, while high performance systems can power their local supplies directly from the output capacitor and run non-critical portions of the chip. Our system uses the capacitor's stored energy to directly power the TEC current source. Once the energy in the capacitor is depleted, we switch to powering the TEC from the chip's voltage supply and ensure that maximum cooling is maintained. Figure 36 illustrates the proposed systems behavior in the time domain. When the chip power is low and the chip temperature is low, we operate in the harvesting mode. The boost regulator boosts the voltage of the TEG and stores the converted thermal energy into an output capacitor. If a high power event occurs in the system, such as a high-power application running, the chip temperature would begin to rise and could go above the thermal limit of the chip and package. This is the common case of turbo-boosting which was discussed previously. With a TEM within the system, once a pre-defined threshold is reached, we can put the system into the cooling mode, reducing the temperature increase rate of the chip and hopefully keeping the temperature below the thermal limit and completing the high-power workload. It is important to note that if the chip power is very high or the TEC current is not large enough, we might not be able to avoid the thermal limit with only the use of the TEC, and might require throttling, thread migration, or a more powerful cooling solution. The TEC will still allow us to extend the workload time.

5.3 System Level Circuit Design
5.3.1 TEC Mode Controller

The TEC Mode Controller is responsible for sensing temperature, and making a decision for what mode the system should operate in, TEC or TEG. The goal of the controller is to reduce the temperature of the silicon chip and avoid the thermal limit of the package, while harvesting maximum energy. For the simplicity of design, the controller senses when the temperature crosses an externally adjustable threshold, and turns the TEC on by pushing a constant current through it. If the chip temperature is below the reference we operate in
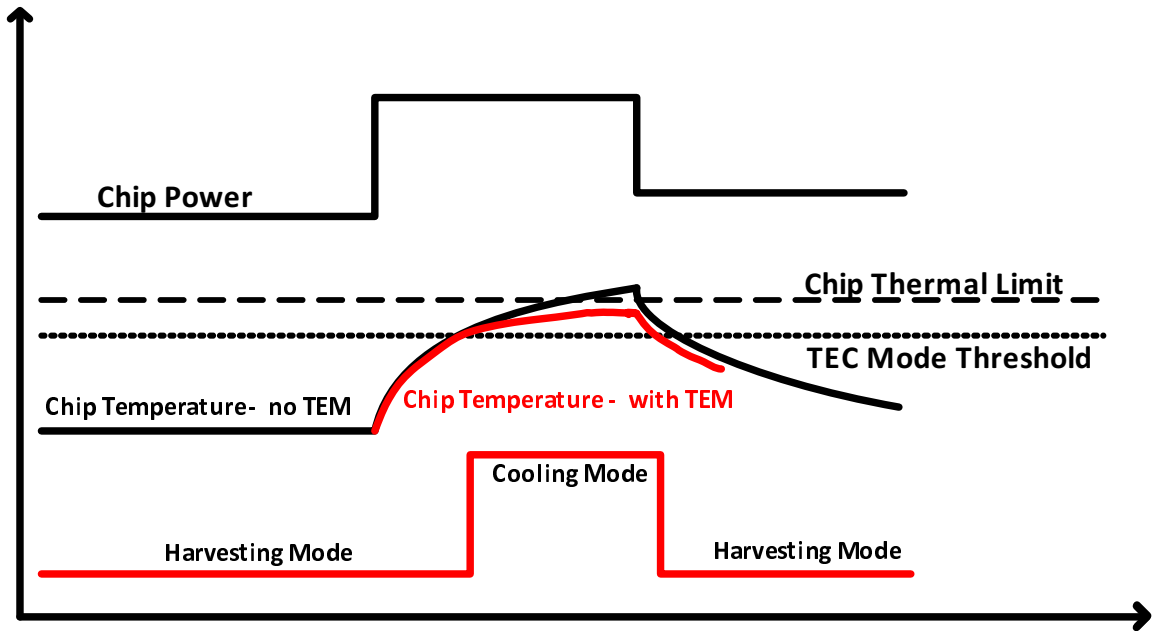
Figure 36: The system level operation of the proposed system.

the TEG mode. Figure 38 shows the top level structure of the mode controller. We have a central temperature sensor that generates 2 analog output representations of temperature. The temperature sensor is implemented on chip using a lateral BJT with a constant current being pushed into the BJT, generating a voltage that varies with temperature, V_{OUT}. The output has inverse temperature dependence and only varies by about 2 mV/°C, which is fairly small. We have therefore added another amplification stage on chip in order to boost this to about 4 mV/°C in order to avoid false triggering of the comparators. This is the V_{AMP} output as shown in figure 38.

A bank of hysteretic comparators, with an externally selectable hysteresis window, is used as the decision engine. We use a 3-bit off-chip select signal for the comparators, with 4 comparators set for each of the temperature outputs. It is important to note that the hysteresis is added to avoid unnecessary mode switching of the comparator when the temperature is hovering near the reference. If the temperature crosses the reference, the comparator trips and the digital signal turns the current source on, pushing current through the TEC
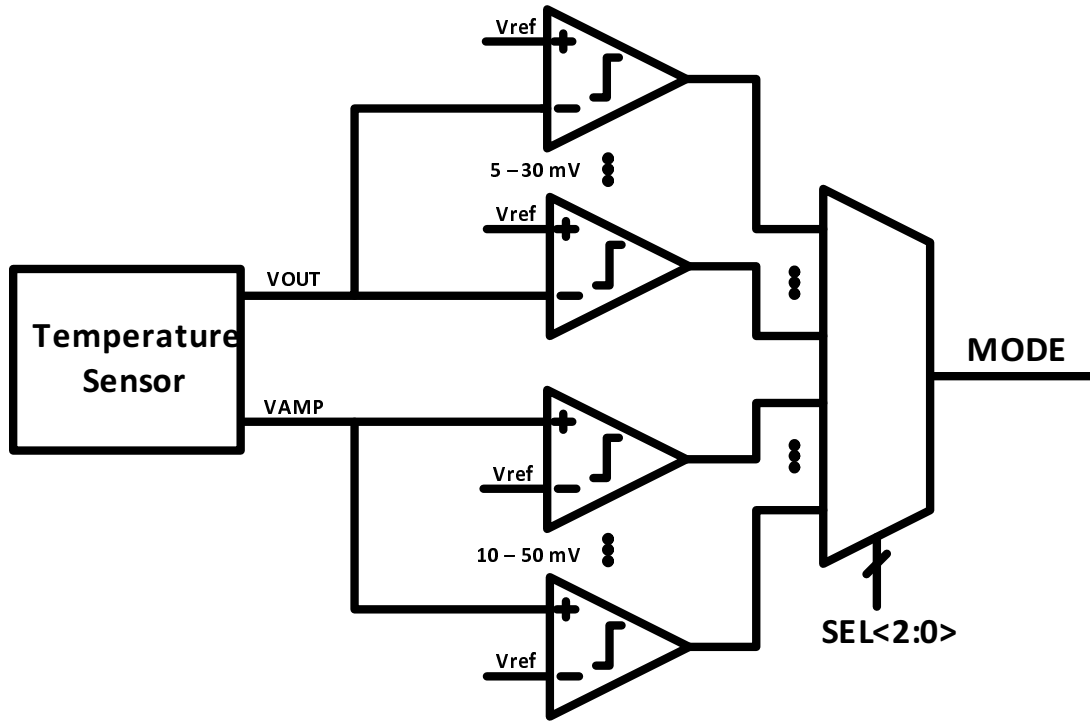
Figure 37: Top level design of the TEC mode controller. The MODE signal drives the switch matrix and puts the system in the appropriate mode.

off-chip. By changing the hysteresis we effectively change the temperature reduction be-fore the comparator switches again and puts the system in the harvesting mode. The digital outputs of the comparator are then fed to a mux that has externally controlled selection, and the output of the mux is buffered to drive the switch matrix. M2 and M3 are turned on during TEC mode.

Figure 39 shows the design of the current source. We use feedback and sense a portion of the output current, divide it down 100X and compare it to the reference current set to make sure that if the source voltage of the output transistor moves, so does the drain in order to keep the current steady. With the I_{SEL} external signal, we can program the reference current from 0-50μA and hence the output current from 10-100mA. This is done by turning the $I_{SEL} < 3:0 >$ switches on and mirroring a higher reference current. Using this feedback approach, the constant current can be maintained across a 10Ω load from 3.3V all the way down to 1.5V, before the feedback cannot respond further. This represents the discharging
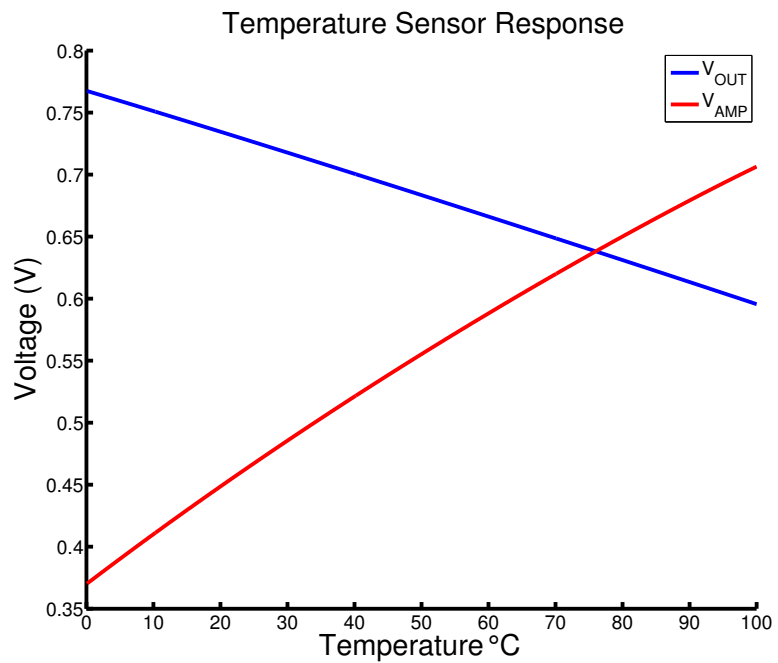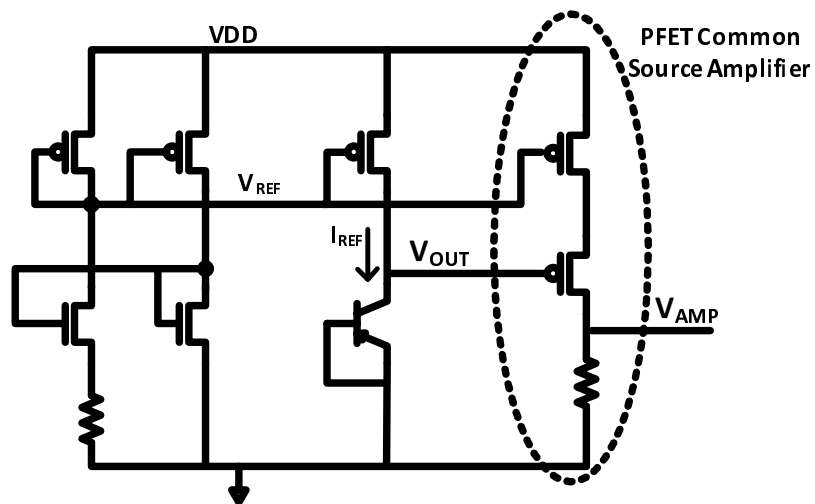
Figure 38: The temperature sensor used on chip: a) schematic design. V_{OUT} is used as the low noise output to detect but since the output variation is small, we have added the V_{AMP} output in order to avoid false triggering in the comparators. b) Simulation results of both outputs.
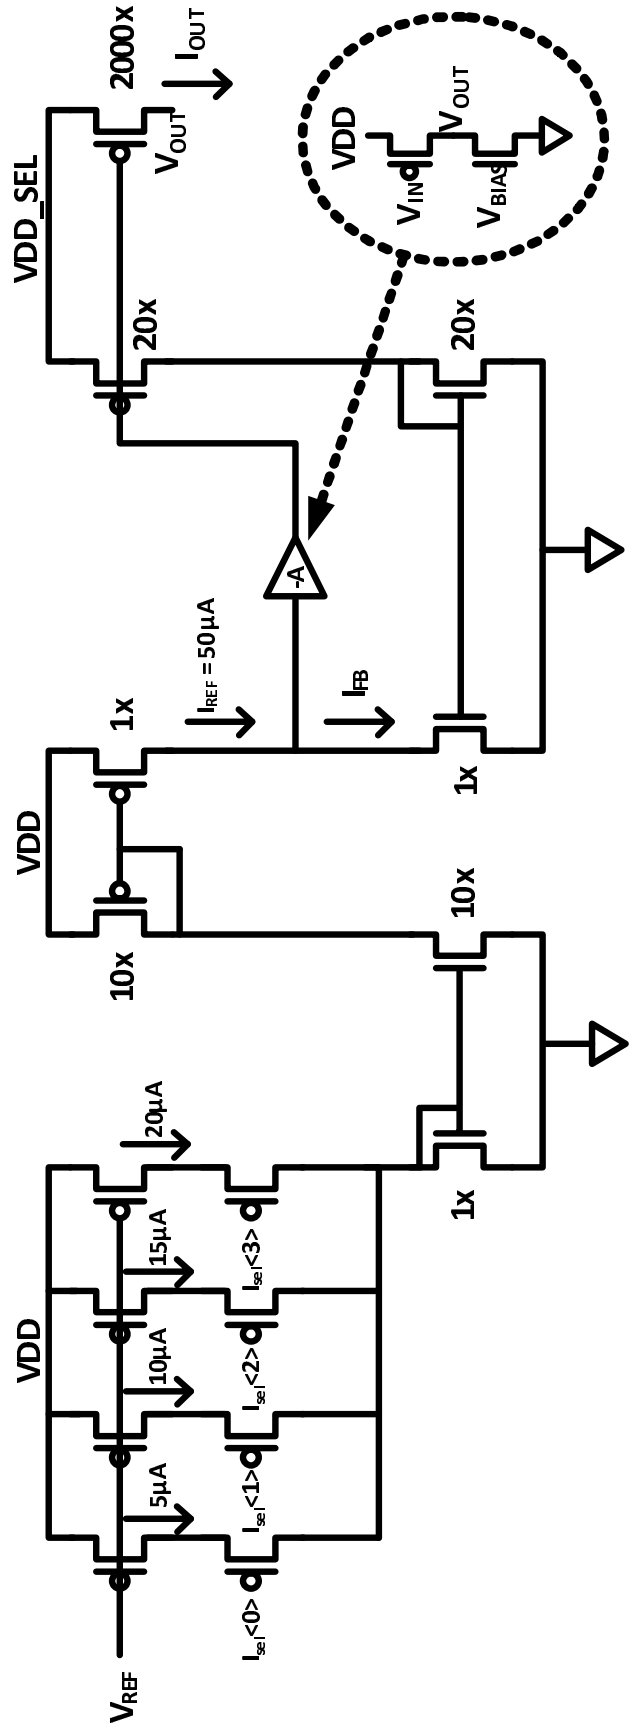
Figure 39: The TEC Current source. $I_{SEL} < 3 : 0 >$ sets the output current externally, and the feedback ensures constant current even with VDD_{SEL} reducing (output capacitor being discharged).

more stored energy from the output capacitor, achieving higher energy efficiency.

5.3.2   TEG Harvesting System
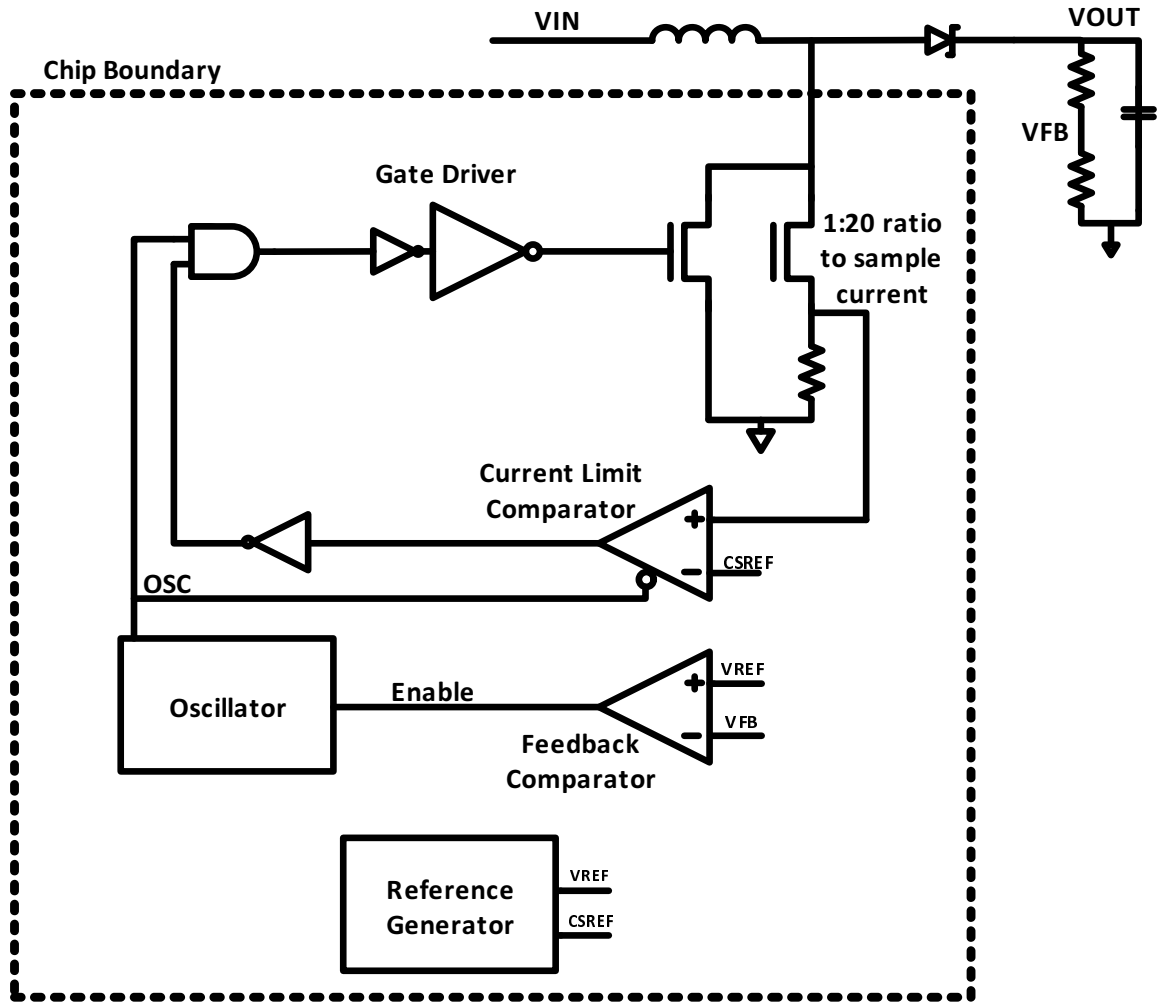
Figure 40: Top level design of the TEG boost regulator. The booster is designed using Pulsed Frequency Modulation (PFM) control and can service various loads.

The test chip also includes an integrated Pulsed Frequency Modulation (PFM) boost regulator to boost the low input voltages that the TEM generates in idle power modes to a high voltage for storage to a capacitor. Figure 40 shows the functional block diagram of the boost regulator. The V_{FB} signal monitors the output voltage and uses a hysteretic comparator to compare it with an internally generated reference. The output of the comparator is low when V_{FB} is higher than V_{REF}, meaning that the output is above the regulation

69

point. This means that the oscillator and current limit comparator are both turned off and the output will discharge by the load condition. Once V_{FB} discharges beyond the reference point, the comparator will flip and enable the oscillator and turn the NFET switch on. This will cause the inductor current to build up linearly while the output continues to discharge and the current buildup can be stopped by the current limit comparator or the end of the oscillator pulse, whichever occurs first. When the peak current is reached, the NFET turns off, and the inductor currents discharges to the load, increasing V_{OUT}. Once the inductor current goes to zero the load will start discharging again, causing the process to restart. If the output load is very large, the output voltage might not increase above the hysteresis and the oscillator will not turn off and continue to pulse. Because the oscillator has a very large duty cycle (98% on time), multiple pulses will continue to build up the current and cause the output to slowly charge up and flip the comparator, restarting the process described above. This allows for our booster to service various loads and provide a high conversion ratio.

5.4 Test Chip Implementation

The proposed design is implemented in an IBM 130nm CMOS process and fabricated through MOSIS. The chip die photo, test board, and external TEM are shown in figure 41. An overall chip summary and measurement characteristics are presented in table 7 The chip has a total die area of 1mm X 1mm with the boost regulator placed at the bottom of the chip and the cooling mode circuitry and decision making circuits in the top left and middle of the chip. The switch matrix is at the top right of the chip and uses large transistors to reduce the losses across the on resistance of the switches. The test chip was supposed to be brought into the clean room to have a TEM directly attached to the backside of the Silicon layer to be able to test the transient cooling developed in this thesis as well as the TEC control methods. The test chip also had an internal resistive heater element, allowing us to control the temperature of the Silicon layer with a maximum heat flux of $100W/cm^2$. This
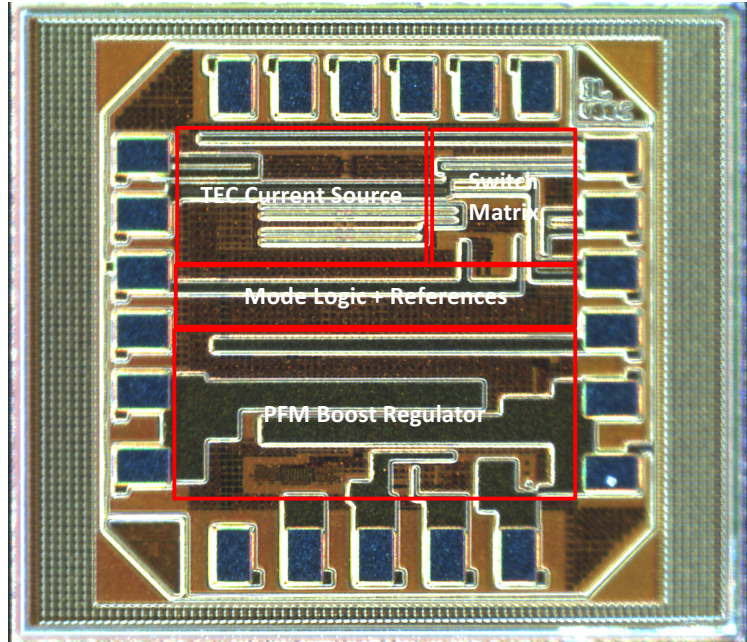
(a)
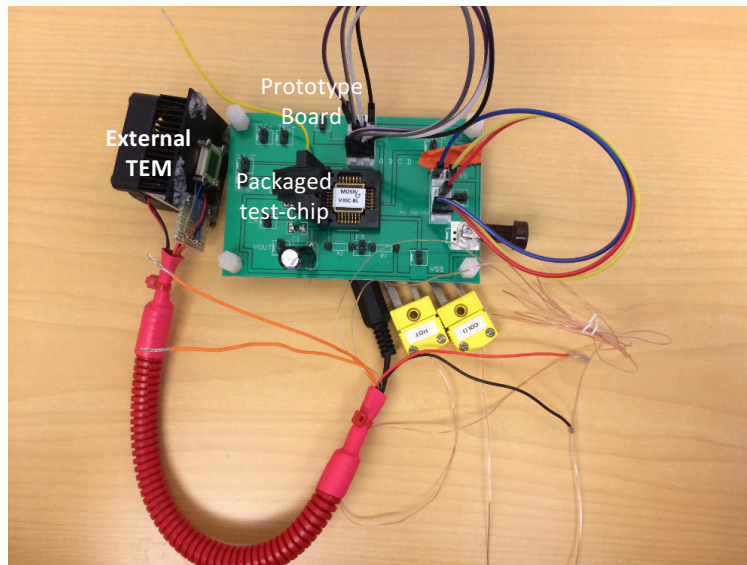
(b)

Figure 41: System implementation: a) test chip and b) test board.

ensured that we could control the temperature on chip and verify the control loop. Unfortunately, due to cost constraints for clean room use as well as a limited number of working dies that came back from MOSIS we were unable to integrate the TEM directly to the Silicon of our chip. Because the TEM could not be directly integrated into the packaging solutions of this test chip, we have characterized our system with an external TEM solution. This external system uses the same TEM element, but has a bulkier thermal solution with a far greater thermal time constant. We use the EV56 TEM Evaluation kit from Laird Technologies (formerly Nextreme) [47]. This TEM has an internal resistance of about 10Ω and a heater and heat sink solution. This serves as the thermal evaluation, with our chip serving as the electrical control system. The thermal heater can apply heat fluxes up to $150W/cm^2$ to the hot side of the TEM while the fan cools the cold side to generate large ΔT. There are 2 thermocouples to monitor the temperature of both the hot and cold side of the TEM. Although the thermal response of the external system is much slower than an integrated device will be, the functionality of the energy management solution is presented and is an ideal candidate for integrated TEM materials. With the ability to function with a 3mmX3mm external TEM, this energy management solution can be integrated under the TEM with about a 10% area overhead on chip. Scaling it down to lower technology nodes can lead to even greater area reductions.

5.5 Measurement Results

5.5.1 TEC Mode Control Characterization

The TEC Mode Controller has been fully characterized and tested using the external TEM. We connect the thermocouple leads to an external instrumentation amplifier (AD8495) [48], in order to read the temperature with accuracy of $0.1°C$ at $5mV/°C$. This ensures an accurate and low noise temperature measurement. We apply a high heat flux to the heater and TEM, $50 \ W/cm^2$ and monitor the temperature at the hot side. We connect the TEM to the test chip and force it in the TEG mode initially using the external voltage (temperature) reference. Figure 42 shows the response of the TEM module. The temperature rises initially as shown

by the yellow line in the scope capture. This is the voltage output of the AD8595 amplifier and has an output of 5mV/°C so we see the temperature increases from 25°C to about 75°C. Once the temperature reaches that level, we force the mode controller into the TEC mode using the reference and turn the TEC current source on with a current of 57mA. We observe a reduction of the hot side temperature of the TEM (equivalently the chip temperature in the package integrated TEM). This results in a steady state temperature reduction of about 5°C, a significant reduction. This result illustrates the concept of using the TEC for temperature reduction of the chip, and avoiding the package temperature limit. The blue and green lines show the positive and negative terminals of the TEC respectively. The current levels of the TEC Current source can be programmed digitally with the external 4-bit I_{SEL} signal as described in section 5.3.1. The measured test chip can successfully source anywhere from 17-105mA, matching closely with the simulated design. Table 8 This allows the chip to be programmed depending on the TEC solution that is attached. With larger TECs we require a larger current to achieve a higher ΔT, but smaller form factor TEMs such as ones to be

Table 7: Test chip description and key measurement parameters

Test Chip Summary	
Technology	IBM 130nm CMOS
Die Size	1mm X1mm
Total System Area	0.49 mm^2
Maximum TEC Current	105 mA
Maximum Boost Regulator Output	3.3 V
Logic Supply Voltage	1.2 V
Total Chip Power Dissipation	5.6 mW
Total VDD Power	0.62 mW
TEC Controller Power	4.98 mW
TEC Current Source Start-up	3.785 ms
TEM Device	Laird Tech EV56
Maximum TEM Heat Flux	150 W/cm^2
TEM Electrical Resistance	10 Ω
Inductor (External)	100 μH
Output Capacitor	1 mF

integrated into mobile devices will require a lower current through the TEC for optimal

cooling. This allows for our solution to be used with different sizes of TEM materials.
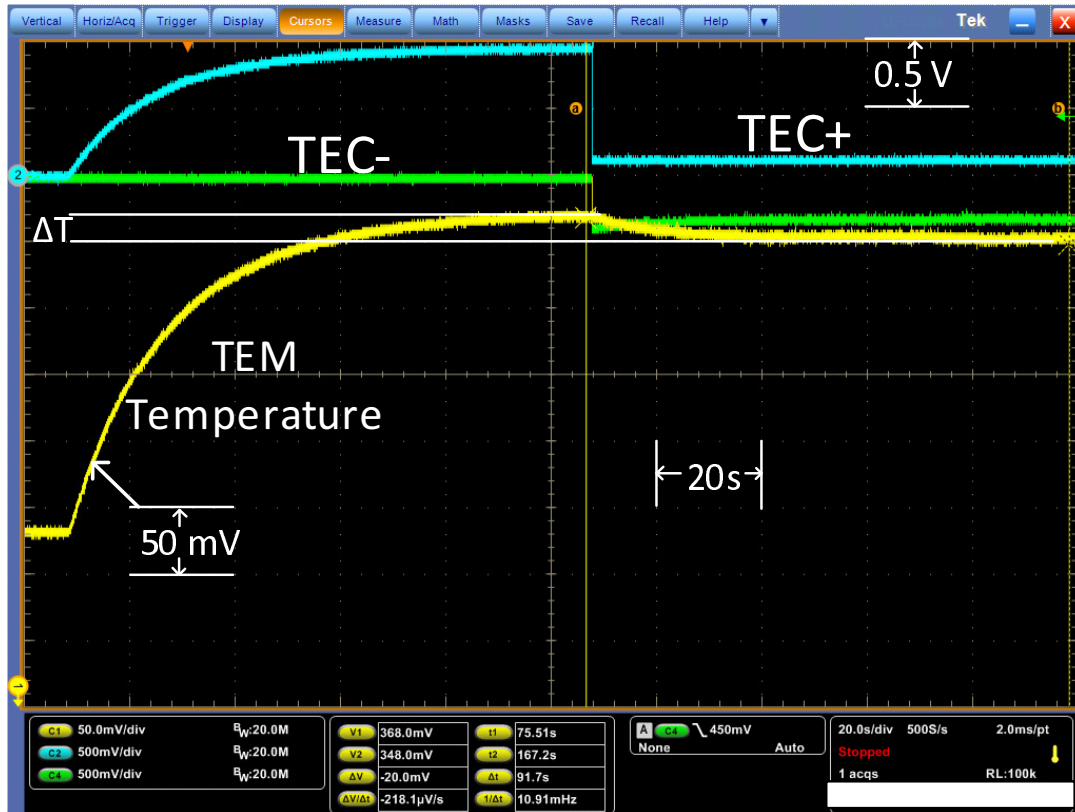
Figure 42: Temperature Characteristics of test chip and external TEM. (a) steady state response and (b) transient response.

Table 8: Current Source Characterization. Only a subset of the 16 setting are shown. Each digital setting is controlled by an off-chip signal and can be changed very easily by a designer to ensure maximum TEC performance in cooling solution.

Digitally controlled current source	
$I_{SEL} < 3 : 0 >$	Output Current
1111	12 μA
1110	17 mA
1101	31 mA
1010	56 mA
1000	78 mA
0011	84 mA
0000	105 mA

Figure 43 shows the maximum attainable ΔT at the hot side for given TEC currents and heat fluxes. We see that up to 105mA the TEM is able to reduce the temperature of the hot side, meaning that for this TEC we prefer to always run maximum current. It is important to note that with smaller form factor TECs the optimal current can change to a lower value, and our chip can have the current source externally programmed for maximum performance of a specific current. This allows us to program a chip externally for it's optimal TEC current. Additionally this can be implemented as a LUT on chip and directly programmed by the chip. This allows for co-design with the architecture. We also observe that we get a larger ΔT for higher heat fluxes, but it is important to note with $100\text{W}/\text{cm}^2$ applied, the initial temperature before turning the TEC on was $108°C$, which might is above the thermal limit for many systems. In the $50\text{W}/\text{cm}^2$ case the initial temperature before the TEC was turned on was $76°C$, making the $7°C$ ΔT more effective from a fraction of the initial temperature standpoint. We also consider the event of the TEC turning on at a pre-defined temperature. Previous work in the literature has investigated this method in order to extend the workload time before a thermal limit [7]. Figure 44 shows the chip and system coupled performance with this method. We turn the TEC on with 57mA at a threshold temperature and observe an immediate cooling effect. We see that the rate of increase in the temperature is decreased and the chip will reach the steady state temperature later. This time allows a chip to complete the workload before having to throttle down.

Lastly it is critical to note the power dissipation of the controller, as well as the time it takes to start-up the current source. Because the controller is based on mostly analog circuits, the major power consumption comes from the biasing networks, and the switching power of the switch matrix transistors. The average power dissipation of the controller, excluding the actual current flowing through the TEC, is 4.98mW. This fairly significant power is due to the feedback network in the current source as noted in section 4.1. Since we need to sample a portion of the output current in order to adjust the voltage, we have to always dissipate a significant biasing current in the current source. We sample 1/100
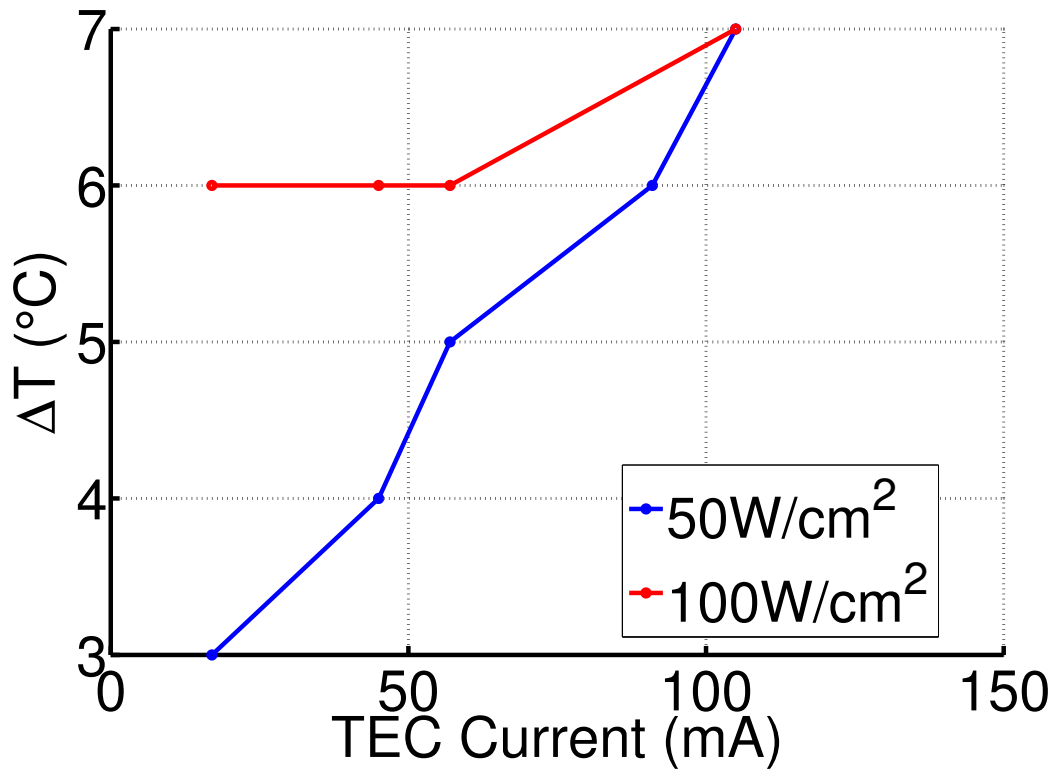
Figure 43: Steady-state temperature reduction for various cooling current levels from the programmable current source and TEM solution.

of the output current, as anything smaller than this leads to large errors in current due to process variation. Despite this, the power is a small portion of the total cooling power required by the system, 346 mW with the maximum current programmed. In addition to the power dissipation it is important to consider the start-up time of the current source, once the decision is made. The TEC current source had a measured start-up time of 3.785ms. With the external system that has cooling times on the order of 10s of seconds this is insignificant, and even with an integrated system, where the cooling time scales are on the order of 100s of ms this controller still responds very quickly.

5.5.2 TEG Harvesting Mode Characterization

The on-chip asynchronous boost regulator has also been characterized for functionality, speed, and efficiency. Figure 45 shows a scope capture showing the functionality of the booster. We use the Nextreme kit again and apply a heat flux to the TEM to generate a

voltage across it. We connect the switch matrix in the booster mode which connects the positive TEM terminal to the inductor. As we can see the output voltage (yellow line) begins to rise in a linear fashion as it is heavily loaded (1mF output capacitor) and the oscillator is operating at the highest switching frequency, measured experimentally to be 93.9kHz. The large output capacitor is used in order to be able to store sufficient energy for cooling. The green curve shows the inductor node that connects to the on-chip FET. As we can see there is continuous switching at the node with the voltage steadily rising. The purple curve is showing the positive TEM terminal voltage. As we can see, initially the TEM is loaded by a large inductor current and hence the voltage is effectively reduced due to the load resistance as well as trace and inductor resistances. As the output builds
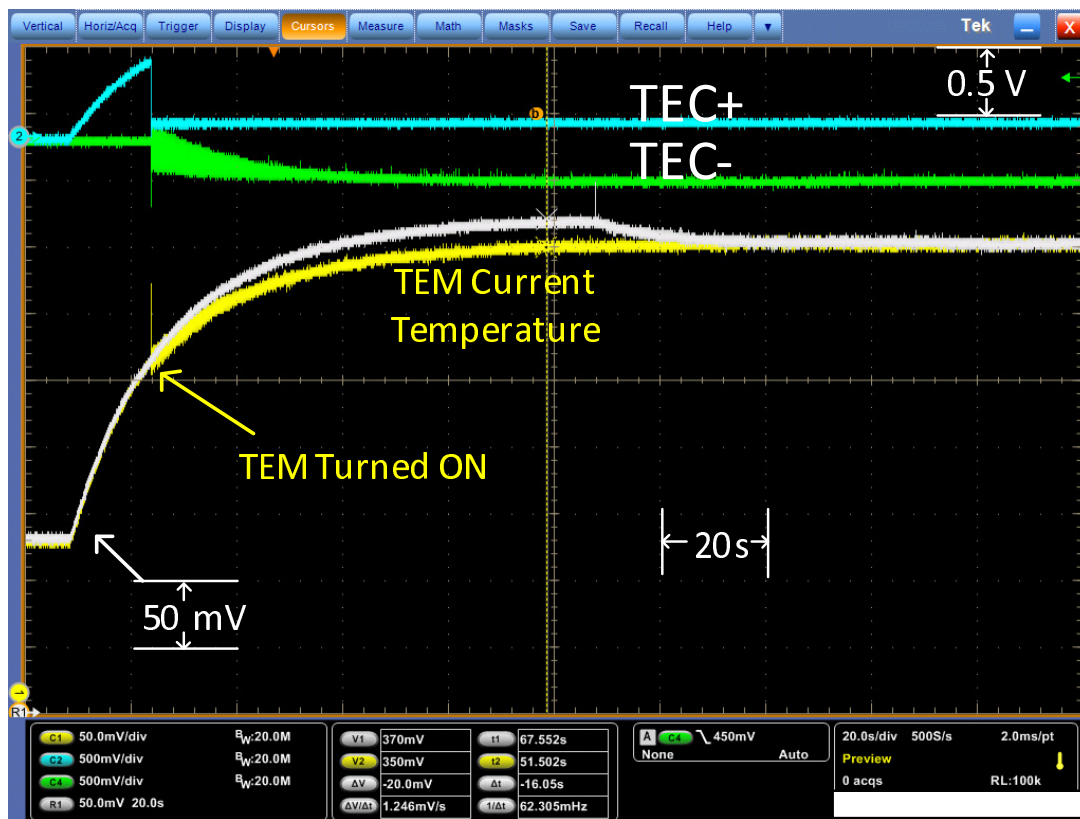
Figure 44: Measured transient temperature characteristics of test chip and external TEM. The yellow curve shows the temperature profile for transient cooling while the white curve is a superimposed steady-state behavior. We observe that the TEC provides instantaneous cooling and can help avoid thermal limits from the chip.
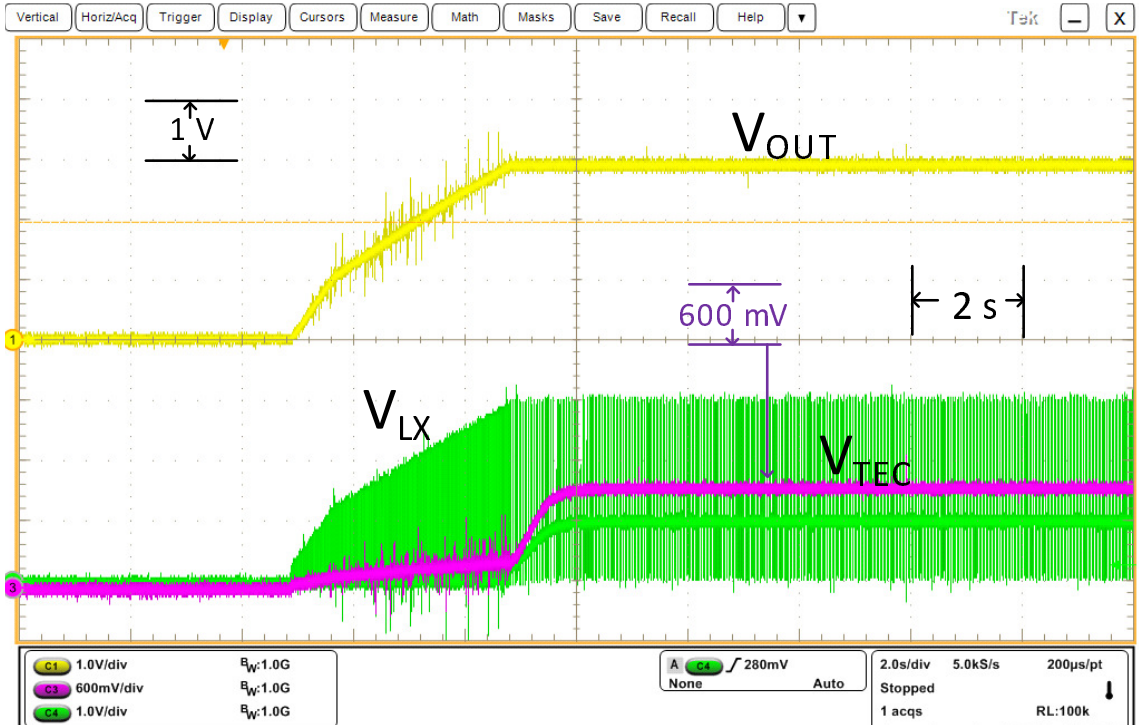
Figure 45: Start-up of PFM boost regulator. The output is regulated at 3V with a 1mF output capacitor.

up and the current drawn reduces, the voltage begins to steadily rise. Once the regulation point is reached the input voltage settles as the loading decreases, and the switching at the inductor node is also reduced. Since the output voltage is quite high, we are still operating near the maximum switching frequency of the oscillator to keep the output regulated at 3V. We have also characterized the efficiency of the booster considering a DC load. As we can see in figure 46 the maximum efficiency of the booster is near 80% at higher output loads. As the load decreases so does the efficiency.

5.5.3 Full System Characterization

Finally the full autonomous system has also been characterized with the external TEM. Figure 47 shows the scope capture of the integrated system. The yellow line is the output voltage and the purple is the input to the boost regulator. As we can see we turn the system on in the TEG mode and begin harvesting energy into the output. As stated in the previous section, the system takes time to boost up as the input voltage is around 900mV (ΔT of
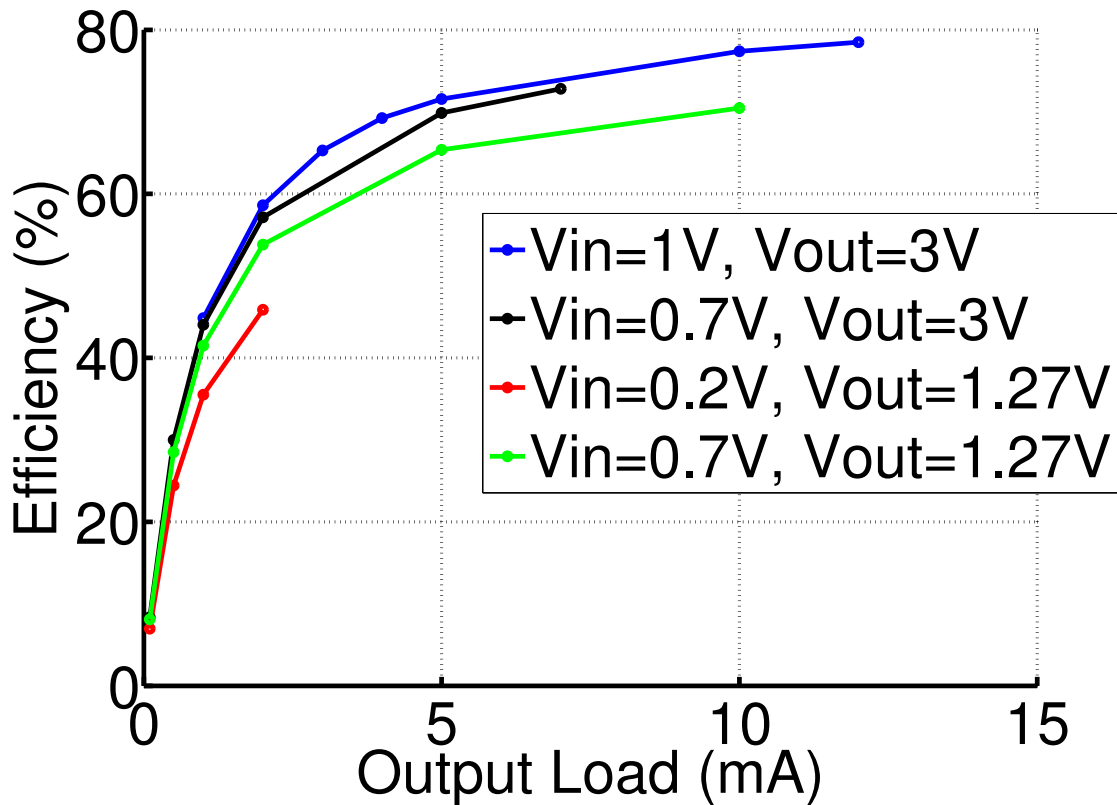
Figure 46: Fixed load efficiency contours of the boost regulator. The booster reaches maximum efficiency near 80%

about 30°C) and the output capacitor is 1mF. Once the system has the output regulated, we force the controller into the TEC mode and we force the TEC mode controllers current source to be powered by the output voltage directly. We can see that the output voltage is quickly discharged by the large current that the TEC requires to cool. As we see from the inset scope capture the discharge is linear, meaning that the current source works despite the reduction in output voltage, and sources a constant current, 37mA as programmed in this case. Before the output is fully discharged, a VDD selection comparator flips and the TEC current continues to be sourced from the chips supply, continuing the cooling operation. The current can be fully sourced from the output capacitor for about 35ms, a time that is quite significant at microprocessor time scales. As the system is fully functional with an external TEM, this low overhead design can be integrated in a chip with an integrated TEM

in the package.
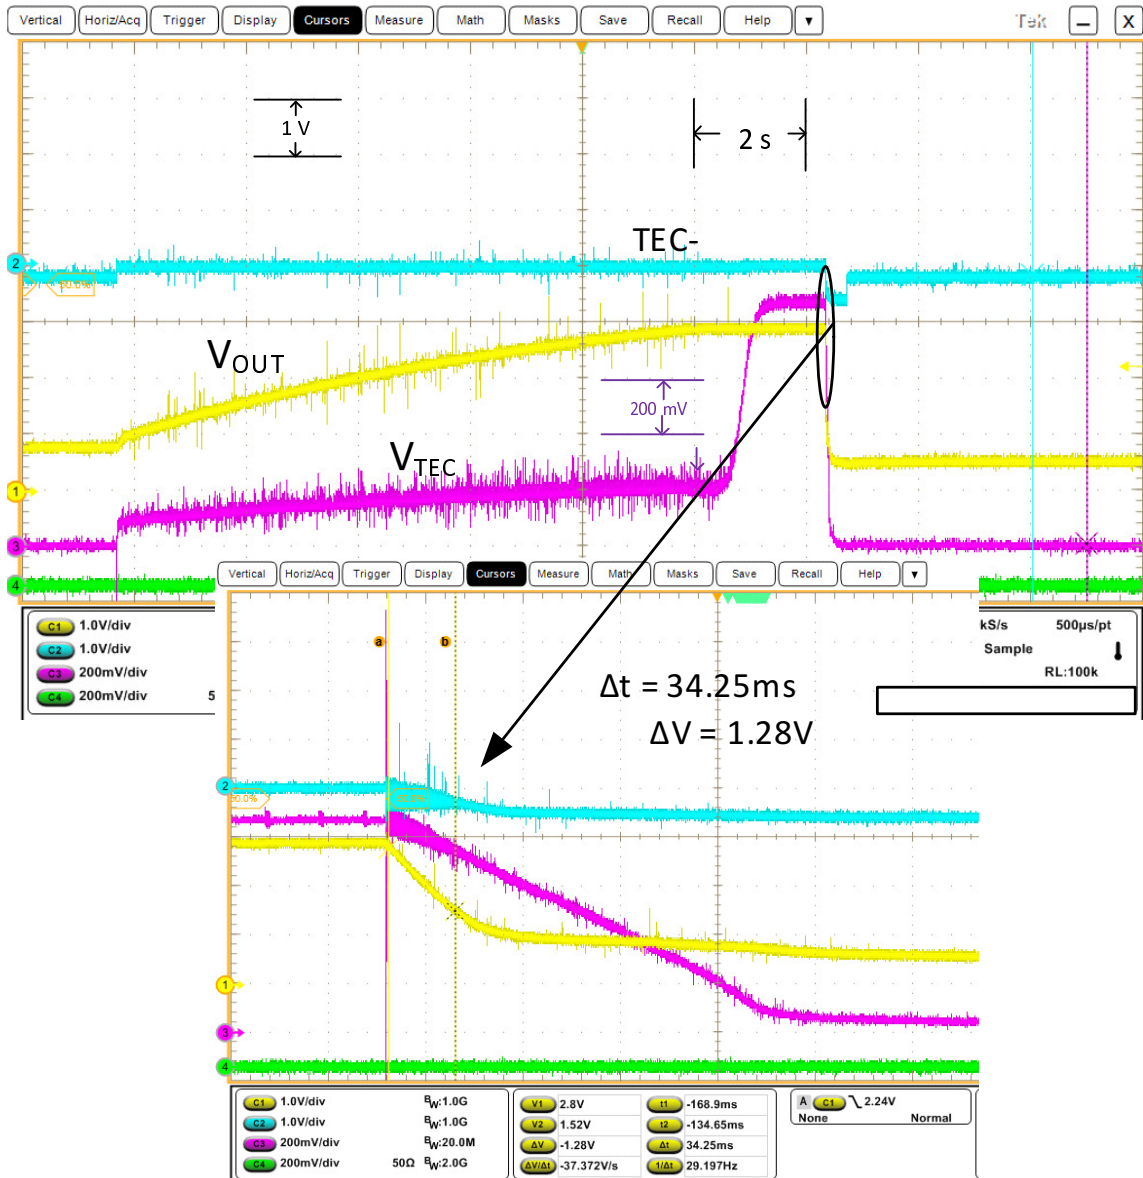
Figure 47: Switching response of full system. The system harvests energy and regulates the output to 3V, before the TEC is turned on the output is used to supply the constant TEC current. This lasts for 34ms when the current source switches to the chip VDD to source the current.

5.6 Conclusion

We have presented the design of an autonomous system that is able to use a single TEM for cooling as well as harvesting energy. Our system is able to harvest energy and store it in an

output capacitor, and when cooling is required draw energy from the output capacitor. This system only takes up 0.49mm2 in 130nm CMOS, consumes 5mW of power (excluding cooling power), and can start-up the cooling operation in 4ms. The system is suitable for chip-scale on-demand cooling. The small area and moderate power make the design suitable for integration in high-performance microprocessors. The integration of the proposed system in microprocessor packages with embedded thermoelectric modules can improve the energy-efficiency of on-demand active cooling. The future design needs to reduce the minimum input voltage for the successful start-up of the booster to allow harvesting from even lower chip power. The response time of the booster needs to be improved as well. The future research in this direction needs to consider demonstration of integrated TEM with the test-chip (instead of external TEM), and the co-design of integrated TEM control and processor architecture to better exploit the energy-efficient on-demand cooling.

CHAPTER 6

CONCLUSIONS

6.1   Contribution

The purpose of this thesis is to develop a methodology for designing a high performance VLSI system using integrated super-lattice based thin-film thermoelectric coolers (TECs) to improve the thermal characteristics, performance, and energy efficiency. Modern VLSI systems have reached the power wall and hence extra performance gains are becoming limited due to the inability to pump out the heat generated by the high density integration of billions of transistors. Therefore, developing a more cost effective cooling solution is critical to extending performance in future high performance systems. This thesis studies the potential of embedding TECs within a package and using the device to manage thermal events in a system. These small form factor devices will not add to system volume and are more importantly active devices. This thesis thoroughly investigates the system level implications of using TECs by developing and verifying a compact thermal model. We use the model to assess the temperature reduction and performance benefits that the TEC contributes to a system in a single and multi-core scenario. We develop control methods to efficiently control the TEC and finally we propose a system to increase the energy efficiency of the overall system and harvest energy during thermally non-critical events. The contributions of the thesis are as follows:

1. The thesis develops the system level thermal model for the TEC and package and thoroughly studies the system level prospects of embedding a TEC within a processor package and evaluating the benefits at the single and multi-core levels.

2. The thesis introduces control methods to aid in the prevention of thermal violations for unknown power profiles. The control methods are implemented within a co-simulation environment that simultaneously simulates electrical and thermal systems.

3. The thesis implements an energy autonomous system that is capable of increasing system efficiency by using the same device for both cooling and harvesting and using harvested energy to cool the chip. The system is implemented in an IBM 130nm process and characterized with a commercial TEM.

6.2 Recommendations for Extension and Future Work

The research in this thesis has assessed the benefits of embedded thermoelectrics but there are many avenues for continuing the research further. As a first step additional studies are required on the architectural level. This thesis has presented the potential to reduce thermal violations and extend performance, but a more thorough study is required that will couple the architecture with the thermal simulation framework so that the exact performance benefits can be seen. This means that a tool needs to be developed that will be able to couple microarchitecture development with the transient thermal simulator so that a designer can quickly run benchmarks for execution time and thermal performance in a magnitude of layouts. Tools like Hotspot [49, 50] are great for evaluation of the maximum temperature of a package with a specific architecture, they fail to account for the transient performance and are not sufficient for modeling DVFS and other power boosting techniques. This will require a thorough investigation into architectural methods to take advantage of transient benefits of TECs while also allowing rapid evaluation of workloads within a newly proposed architecture.

The control methods presented in this work only scratch the surface of efficient TEC control. More work can be done to develop more advanced control methods that optimize across various design parameters and achieve optimal chip temperature with minimum power. These controllers should have a rapid design flow and should be implemented with low overhead on chip. The flow should be implemented within the standard design methodology of IC development and allow technology/architecture/circuit interactions. Although the work here has characterized the TEC control on chip, future work in this area

83

should attach the TEM to a chip and evaluate the performance benefits further and take advantage of the transient cooling effects. Attaching the TEC to the chip and connecting the terminals to the Silicon requires packaging innovations and should be investigated further. TEC system benefits have also been studied for 3D systems where thermals are more difficult to control due to many layers blocking the heat flow. The controls methods should be extended further to account for 3D systems and evaluate the efficacy. Mobile systems should also be studied within the thermal modeling framework so that the TEC effects can be understood. Mobile systems have even more stringent temperature limits due to low skin temperature limits and varying workloads that can quickly cause thermal throttling. Investigating the TEC performance benefits is critical to assess the feasibility of using them within a mobile system, where thickness and space requirements can be very limited.

Finally the autonomous system proposed here needs to be investigated further. Since the TEM is often optimized for either cooling or harvesting, a package should integrate materials optimized for both modes and placed within the chip layout. This would allow the system to harvest a greater amount of energy while providing maximum cooling in hot-spot locations. Additionally the autonomous system's chip implementation should have a TEM directly attached using advanced packaging techniques.

6.3   Critical Assessment

No thesis would be complete without the author's critical assessment of the work. The research undertaken has shown the potential that the TEC can have on thermally constrained systems, by allowing for short bursts of cooling and extended turbo-boosting. Although the simulations results are promising, the TEC was not integrated with a processor and real workload so that real lab measurements could be verified. In addition, the TEC device physics have significant effect on the cooling performance, and TEC device optimizations are required to be able to optimize the device to provide maximum cooling capability at

the lowest power budget. This requires device engineering and fabrication techniques to be developed further so that a system designer has the ability to use the most effective device in the system. In addition, manufacturing of the TEC device at the back of the heat spreader requires significant processing and cost, which might render the TEC device cost prohibitive for a chip manufacturer to integrate. If we can show that using an integrated TEC is the only way for thermal mitigation in the future, this might reduce the overall cost of manufacturing due to scale and make the integrated TEC part of future chip generations.

The thesis did not critically benchmark the TEC to other available cooling technologies. This is required if this technology is going to be developed further and integrated onto next generation chips. Criteria for comparison needs to be developed so that competing cooling technologies can be compared from an overall system area and cost perspective so the most efficient and economical system can be chosen. It might not make sense for a system to include a TEC if simply adding a more powerful fan can fulfill performance benefits.

Lastly it is critical to note that the TEG energy harvesting has a poor efficiency, less than 1%. This requires significant improvements in the device engineering to improve the conversion efficiency of the material. Again this needs to be evaluated further and compared to competing technologies so that the system has the highest efficiency overall without prohibitive cost.

REFERENCES

[1] E. Rotem, A. Naveh, D. Rajwan, A. Ananthakrishnan, and E. Weissmann, "Power-management architecture of the intel microarchitecture code-named sandy bridge," *Micro, IEEE*, vol. 32, pp. 20–27, March 2012.

[2] R. Dennard, F. Gaensslen, V. Rideout, E. Bassous, and A. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," *Solid-State Circuits, IEEE Journal of*, vol. 9, pp. 256–268, Oct 1974.

[3] R. Venkatasubramanian, E. Siivola, T. Colpitts, and B. O'quinn, "Thin-film thermoelectric devices with high room-temperature figures of merit," *Nature*, vol. 413, no. 6856, pp. 597–602, 2001.

[4] S. Borkar, "Design challenges of technology scaling," *Micro, IEEE*, vol. 19, pp. 23–29, Jul 1999.

[5] R. Gonzalez, B. Gordon, and M. Horowitz, "Supply and threshold voltage scaling for low power CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 32, pp. 1210–1216, Aug 1997.

[6] S.-W. Sun and P. Tsui, "Limitation of CMOS supply-voltage scaling by MOSFET threshold-voltage variation," *Solid-State Circuits, IEEE Journal of*, vol. 30, pp. 947–949, Aug 1995.

[7] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand, "Leakage current mechanisms and leakage reduction techniques in deep-submicrometer CMOS circuits," *Proceedings of the IEEE*, vol. 91, pp. 305–327, Feb 2003.

[8] C. Auth *et al.*, "A 22nm high performance and low-power CMOS technology featuring fully-depleted tri-gate transistors, self-aligned contacts and high density MIM capacitors," in *VLSI Technology (VLSIT), 2012 Symposium on*, pp. 131–132, June 2012.

[9] N. Kurd, M. Chowdhury, E. Burton, T. Thomas, C. Mozak, B. Boswell, M. Lal, A. Deval, J. Douglas, M. Elassal, A. Nalamalpu, T. Wilson, M. Merten, S. Chennupaty, W. Gomes, and R. Kumar, "5.9 Haswell: A family of IA 22nm processors," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2014 IEEE International*, pp. 112–113, Feb 2014.

[10] A. Watwe and R. Viswanath, "Thermal implications of non-uniform die power map and CPU performance," in *Proceedings of InterPACK*, vol. 3, pp. 6–11, 2003.

[11] R. Mahajan, C. pin Chiu, and G. Chrysler, "Cooling a Microprocessor Chip," *Proceedings of the IEEE*, vol. 94, pp. 1476–1486, Aug 2006.

[12] R. Kumar, K. Farkas, N. Jouppi, P. Ranganathan, and D. Tullsen, "Single-ISA heterogeneous multi-core architectures: the potential for processor power reduction," in *Microarchitecture, 2003. MICRO-36. Proceedings. 36th Annual IEEE/ACM International Symposium on*, pp. 81–92, Dec 2003.

[13] P. Chaparro, J. Gonzalez, G. Magklis, C. Qiong, and A. Gonzalez, "Understanding the Thermal Implications of Multi-Core Architectures," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 18, pp. 1055–1065, Aug 2007.

[14] J. Donald and M. Martonosi, "Techniques for Multicore Thermal Management: Classification and New Exploration," in *Computer Architecture, 2006. ISCA '06. 33rd International Symposium on*, pp. 78–88, 2006.

[15] Y. Ge, P. Malani, and Q. Qiu, "Distributed task migration for thermal management in many-core systems," in *Design Automation Conference (DAC), 2010 47th ACM/IEEE*, pp. 579–584, June 2010.

[16] P. Pillai and K. G. Shin, "Real-time Dynamic Voltage Scaling for Low-power Embedded Operating Systems," in *Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles*, SOSP '01, (New York, NY, USA), pp. 89–102, ACM, 2001.

[17] W. Kim, M. Gupta, G.-Y. Wei, and D. Brooks, "System level analysis of fast, per-core DVFS using on-chip switching regulators," in *High Performance Computer Architecture, 2008. HPCA 2008. IEEE 14th International Symposium on*, pp. 123–134, Feb 2008.

[18] R. Wirtz, N. Zheng, and D. Chandra, "Thermal management using "dry" phase change material," in *Semiconductor Thermal Measurement and Management Symposium, 1999. Fifteenth Annual IEEE*, pp. 74–82, March 1999.

[19] R. Kandasamy, X.-Q. Wang, and A. S. Mujumdar, "Transient cooling of electronics using phase change material (PCM)-based heat sinks," *Applied Thermal Engineering*, vol. 28, no. 89, pp. 1047 – 1057, 2008.

[20] M. Bakir, C. King, D. Sekar, H. Thacker, B. Dang, G. Huang, A. Naeemi, and J. Meindl, "3D heterogeneous integrated systems: Liquid cooling, power delivery, and implementation," in *Custom Integrated Circuits Conference, 2008. CICC 2008. IEEE*, pp. 663–670, Sept 2008.

[21] I. Chowdhury, R. Prasher, K. Lofgreen, G. Chrysler, S. Narasimhan, R. Mahajan, D. Koester, R. Alley, and R. Venkatasubramanian, "On-chip cooling by superlattice-based thin-film thermoelectrics," *Nature Nanotechnology*, vol. 4, no. 4, pp. 235–238, 2009.

[22] G. J. Snyder, M. Soto, R. Alley, D. Koester, and B. Conner, "Hot spot cooling using embedded thermoelectric coolers," in *Semiconductor Thermal Measurement and Management Symposium, 2006 IEEE Twenty-Second Annual IEEE*, pp. 135–143, March 2006.

[23] T. Harman, P. Taylor, M. Walsh, and B. LaForge, "Quantum dot superlattice thermo-electric materials and devices," *Science*, vol. 297, no. 5590, pp. 2229–2232, 2002.

[24] D. Mitrani, J. Salazar, A. Turó, M. J. García, and J. A. Chávez, "One-dimensional modeling of TE devices considering temperature-dependent parameters using SPICE," *Microelectronics Journal*, vol. 40, no. 9, pp. 1398–1405, 2009.

[25] D. Mitrani, J. Salazar, A. Turó, M. J. García, and J. A. Chávez, "Transient distributed parameter electrical analogous model of TE devices," *Microelectronics Journal*, vol. 40, no. 9, pp. 1406–1410, 2009.

[26] J. Long and S. Memik, "A framework for optimizing thermoelectric active cooling systems," in *Design Automation Conference (DAC), 2010 47th ACM/IEEE*, pp. 591–596, June 2010.

[27] J. Long, S. Memik, and M. Grayson, "Optimization of an on-chip active cooling system based on thin-film thermoelectric coolers," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2010*, pp. 117–122, March 2010.

[28] J. Long, D. Li, S. Memik, and S. Ulgen, "Theory and Analysis for Optimization of On-Chip Thermoelectric Cooling Systems," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 32, pp. 1628–1632, Oct 2013.

[29] P. Chaparro, J. González, Q. Cai, and G. Chrysler, "Dynamic Thermal Management Using Thin-film Thermoelectric Cooling," in *Proceedings of the 14th ACM/IEEE International Symposium on Low Power Electronics and Design*, ISLPED '09, (New York, NY, USA), pp. 111–116, ACM, 2009.

[30] S. Choday, C. Lu, V. Raghunathan, and K. Roy, "On-chip energy harvesting using thin-film thermoelectric materials," in *Semiconductor Thermal Measurement and Management Symposium (SEMI-THERM), 2013 29th Annual IEEE*, pp. 99–104, March 2013.

[31] T. D. Sands and Y. Wang, "Dynamic switching thermoelectric thermal management systems and methods," Oct. 24 2011. US Patent App. 13/279,475.

[32] S. Choday, M. Lundstrom, and K. Roy, "Prospects of Thin-Film Thermoelectric Devices for Hot-Spot Cooling and On-Chip Energy Harvesting," *Components, Packaging and Manufacturing Technology, IEEE Transactions on*, vol. 3, pp. 2059–2067, Dec 2013.

[33] K. Chang, C. Tsai, C. Teng, and S. Kang, *Electrothermal Analysis of VLSI Systems*. Springer, 2000.

[34] T. Thonhauser, G. D. Mahan, L. Zikatanov, and J. Roe, "Improved supercooling in transient thermoelectrics," *Applied Physics Letters*, vol. 85, no. 15, pp. 3247–3249, 2004.

[35] F. P. Incropera, D. P. Dewitt, T. L. Bergman, and A. S. Lavine, *Fundamentals of Heat and Mass Transfer*. Wiley, 6 ed., 2006.

[36] O. Sullivan, M. P. Gupta, S. Mukhopadhyay, and S. Kumar, "Thermoelectric coolers for thermal gradient management on chip," in *ASME 2010 International Mechanical Engineering Congress and Exposition*, pp. 187–195, American Society of Mechanical Engineers, 2010.

[37] J. Charles, P. Jassi, N. Ananth, A. Sadat, and A. Fedorova, "Evaluation of the intel core-i7 turbo boost feature," in *Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on*, pp. 188–197, Oct 2009.

[38] T. Constantinou, Y. Sazeides, P. Michaud, D. Fetis, and A. Seznec, "Performance implications of single thread migration on a chip multi-core," *SIGARCH Comput. Archit. News*, vol. 33, pp. 80–91, Nov. 2005.

[39] J. Donald and M. Martonosi, "Techniques for multicore thermal management: Classification and new exploration," in *Computer Architecture, 2006. ISCA '06. 33rd International Symposium on*, pp. 78–88, 2006.

[40] K. Woo, S. Meninger, T. Xanthopoulos, E. Crain, D. Ha, and D. Ham, "Dual-DLL-based CMOS all-digital temperature sensor for microprocessor thermal monitoring," in *Solid-State Circuits Conference - Digest of Technical Papers, 2009. ISSCC 2009. IEEE International*, pp. 68–69,69a, Feb 2009.

[41] H. Lakdawala, Y. Li, A. Raychowdhury, G. Taylor, and K. Soumyanath, "A 1.05V 1.6 mW, 0.45°C 3σ Resolution $\Sigma\Delta$ Based Temperature Sensor With Parasitic Resistance Compensation in 32 nm Digital CMOS Process," *Solid-State Circuits, IEEE Journal of*, vol. 44, pp. 3621–3630, Dec 2009.

[42] J. Bierschenk and M. Gilley, "Assessment of tec thermal and reliability requirements for thermoelectrically enhanced heat sinks for cpu cooling applications," in *Thermoelectrics, 2006. ICT '06. 25th International Conference on*, pp. 254–259, Aug 2006.

[43] "AON2400 datasheet." http://aosmd.com/res/data_sheets/AON2400.pdf.

[44] G. Loh, S. Subramaniam, and Y. Xie, "Zesto: A cycle-level simulator for highly detailed microarchitecture exploration," in *Performance Analysis of Systems and Software, 2009. ISPASS 2009. IEEE International Symposium on*, pp. 53–64, April 2009.

[45] S. Li, J.-H. Ahn, R. Strong, J. Brockman, D. Tullsen, and N. Jouppi, "McPAT: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Microarchitecture, 2009. MICRO-42. 42nd Annual IEEE/ACM International Symposium on*, pp. 469–480, Dec 2009.

[46] S. Kanev, G.-Y. Wei, and D. Brooks, "XIOSim: Power-performance Modeling of Mobile x86 Cores," in *Proceedings of the 2012 ACM/IEEE International Symposium on Low Power Electronics and Design*, ISLPED '12, (New York, NY, USA), pp. 267–272, ACM, 2012.

[47] "Laird Tech EV56 datasheet." http://lairdtech.thomasnet.com/item/thermoelectric-modules-2/etec-series/hv56-72-f2-0203-gg.

[48] "Analog Devices 8495 datasheet." http://www.analog.com/en/mems-sensors/digital-temperature-sensors/ad8495/products/product.html.

[49] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. Stan, "Hotspot: a compact thermal modeling methodology for early-stage vlsi design," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. 14, pp. 501–513, May 2006.

[50] W. Huang, *HotSpotA Chip and Package Compact Thermal Modeling Methodology for VLSI Design*. PhD thesis, University of Virginia, 2007.