

COMPOSITE THERMAL CAPACITORS FOR TRANSIENT
THERMAL MANAGEMENT OF MULTICORE MICROPROCESSORS

A Dissertation

Presented to

The Academic Faculty

by

Craig Elkton Green

In Partial Fulfillment

of the Requirements for the Degree

Doctor of Philosophy in the

George W Woodruff School of Mechanical Engineering

Georgia Institute of Technology

August 2012

COMPOSITE THERMAL CAPACITORS FOR TRANSIENT
THERMAL MANAGEMENT OF MULTICORE MICROPROCESSORS

Approved by:

Dr. Andrei G Fedorov, Advisor
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Yogendra K. Joshi, Advisor
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Baratunde Cola
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Muhannad Bakir
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Dr. Sudhakar Yalamanchili
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Date Approved: May 16th, 2012

To Sophia, my parents, Sharon and Herman, and all my loved ones who inspire
and support me.

ACKNOWLEDGEMENTS

I would like to thank my advisors, Dr. Andrei Fedorov and Dr. Yogendra Joshi for their guidance, time, and tutelage throughout my time at Georgia Tech. With your efforts I have become a better engineer and scholar. Likewise, I would like to thank my PhD committee, Dr. Baratunde Cola, Dr. Muhannad Bakir and Dr. Sudhakar Yalamanchili for their time and feedback throughout my PhD proposal and defense process.

I would like to thank the members of my research groups, the Multiscale Integrated Thermofluidics Laboratory and the Microelectronics and Emerging Technology Thermal Laboratory, for their valuable time, feedback and suggestions that have helped to streamline my research process.

I would like to acknowledge the support of the Interconnect Focus Center, one of five research centers funded under the Focus Center Research Program, a Semiconductor Research Corporation program- for both their financial support and technical input that has helped to make my research more meaningful.

Finally I would like to thank my family. My wife, Sophia, I thank for her endless support, patience, encouragement and understanding. My mother I would like to thank for her unconditional love, which has carried me through difficult times. My father I would like to thank for the lessons of hard work, perseverance, and audacity that he taught me as a child and which drive me to continually strive to be the best version of myself.

TABLE OF CONTENTS

| | |
|---|------|
| Acknowledgements | iv |
| List of Tables | viii |
| List of Figures | ix |
| Nomenclature | xiii |
| Summary | xvii |
| Chapter 1: Introduction | 1 |
| 1.1 Introduction | 1 |
| 1.2 Proposed Approach and Relevant Work | 4 |
| Chapter 2: Spreading Impact Analysis | 10 |
| 2.1 Model Description | 10 |
| 2.2 Numerical Modeling | 16 |
| 2.3 Grid and Residual Convergence Investigation | 19 |
| 2.4 Results and Interpretation | 20 |
| Chapter 3: Analysis of Design Space: CTC performance evaluation- 2D Modeling | 26 |
| 3.1 CTC Layout Description: Si Islands | 26 |
| 3.2 CTC Layout Description: Diamond Shaped Loops | 29 |
| 3.3 Model Description | 32 |
| 3.4 Grid and Convergence Study | 33 |
| 3.5 Results | 34 |
| Chapter 4: 3D Model-CTC integrated into a 3 layer stack with SSC regeneration | 40 |
| 4.1 Model Description | 40 |
| 4.2 Model Implementation and boundary conditions | 45 |

| | |
|--|-----|
| 4.3 Grid and Residual Independence Study | 49 |
| 4.4 Results | 51 |
| Chapter 5: Experimental Characterization | 54 |
| 5.1 CTC on Pyrex Test Chip | 55 |
| 5.1.1 Device fabrication and experimental setup | 55 |
| 5.1.2 Experimental procedure | 59 |
| 5.1.3 Investigation of TIM impact on device performance | 61 |
| 5.2 Monolithic Integration in Si | 62 |
| 5.3 Regenerative Cooling | 65 |
| 5.3.1 Experimental setup | 65 |
| 5.3.2 Device operation and characterization process | 73 |
| 5.3.3 Uncertainty Analysis | 76 |
| 5.4 Results | 79 |
| 5.4.1 CTC on Pyrex Test Chip | 79 |
| 5.4.2 Monolithic Integration in Si | 83 |
| 5.4.3 Regenerative Cooling | 88 |
| Chapter 6: Comparison of Model Predictions to Experiment | 96 |
| 6.1 Problem Formulation | 96 |
| 6.2 Results | 99 |
| Chapter 7: Conclusions and Future Work | 104 |
| 7.1 Summary | 104 |
| 7.2 Conclusions and Key Findings | 105 |
| 7.3 Original Contributions | 107 |

| | |
|--|-----|
| Appendix A: Experimental Device Fabrication | 110 |
| A.1 Pyrex Test Chip | 110 |
| A.2 Standalone CTC fabrication | 116 |
| A.3 PCM selection and loading process | 119 |
| A.4 Monolithically integrated CTC and heater/RTD test chip | 122 |
| References | 128 |

LIST OF TABLES

| | Page |
|---|------|
| Table 2.1 Thermophysical properties evaluated for cylindrical spreading investigation | 17 |
| Table 2.2 PCM material properties evaluated for grid and residual convergence of the numerical simulation | 20 |
| Table 3.1 Properties used in simulations for CTC design analysis | 27 |
| Table 3.2 Grid and convergence study for CTC design analysis | 33 |
| Table 4.1 Material properties used in 3D computational analysis | 45 |
| Table 5.1 Pulse length and heat fluxes of the four configurations for CTC regeneration | 74 |
| Table 5.2 Uncertainty analysis | 78 |
| Table A.1 Si_3N_4 deposition process parameters | 114 |
| Table A.2 Si_3N_4 RIE etch process parameters | 114 |
| Table A.3 PT-ICP oxide mask etch parameters | 117 |
| Table A.4 Bosch process parameters | 118 |
| Table A.5 RCA clean parameters | 123 |
| Table A.6 Photoresist bake times for Si substrate | 123 |
| Table A.7 SiO_2 PECVD deposition parameters | 124 |
| Table A.8 Si_3N_4 RIE Etch Process Parameters | 126 |

LIST OF FIGURES

| | |
|--|----|
| Figure 1.1(a) Impact of increased thermal capacitance on core hopping frequency (b) Impact of regenerative cooling on device utilization | 5 |
| Figure 1.2 Schematic of CTC integration in a 3D chip stack | 6 |
| Figure 2.1 Model Schematic for Cylindrical Spreading Analysis | 11 |
| Figure 2.2 Limits on performance modulation of device operating times through spreading into a PCM alone for $h_{sl}=100$ kJ/kg and $c_p=1$ kJ/kg-°C | 22 |
| Figure 2.3 Impact of effective thermal conductivity (κ) and density (ρ) on cylindrical spreading temperature histories: $h_{sl}=100$ kJ/kg and $c_p=1$ kJ/kg-°C | 23 |
| Figure 2.4 Impact of St on Cylindrical Spreading Temperature Histories for $\kappa=10$ W/m-°C, $\rho=10,000$ kg/m ³ | 24 |
| Figure 3.1 Schematic of Si-Islands layouts (a) isotropic array (b) anisotropic array- all three island sizes shown (dimensions in m) | 28 |
| Figure 3.2 Schematic of 2D computational domain for CTC modeling- Diamond shaped loops | 31 |
| Figure 3.3 Time-on-a-core before reaching a 90°C threshold for the isotropic array of Si-islands layout, and different material properties (for quantitative comparison, a baseline of Si only layer yields 3 ms time of operation before reaching a 90°C threshold) | 35 |
| Figure 3.4 Melt front location at throttle time for Si-Islands layout | 36 |
| Figure 3.5 Time-on-a-core before reaching a 90°C threshold for the anisotropic array of Si-islands layout and different material properties | 37 |

| | |
|--|----|
| Figure 3.6 Time-on-a-core before reaching a 90°C threshold for the diamond shaped loop layout and different material properties | 39 |
| Figure 4.1 Schematic of 3D computational domain and boundary conditions | 41 |
| Figure 4.2 SSC integration approaches (a) SSC integrated into vertically adjacent tier (as modeled) (b) SSC is in same tier as CTC | 43 |
| Figure 4.3 Grid and convergence study for 3D computational domain | 50 |
| Figure 4.4 Results of 3D numerical simulations for CTC performance with and without SSC regeneration: (a) 500 W/cm^2 hotspot heat flux (b) 1 kW/cm^2 hotspot heat flux | 52 |
| Figure 5.1 Schematics of the experimental setups (a) CTC coupled via TIM to Pyrex test chip (b) Monolithic integration of CTC and test chip in Si | 54 |
| Figure 5.2 (a) Hotspot heater layer and RTD layout (b) CTC filled with PCM used in experiments | 56 |
| Figure 5.3 Sample Pt heater calibration curve | 57 |
| Figure 5.4 Partially filled CTC, after Pt coating for improved wettability | 58 |
| Figure 5.5 Electric diagram of CTC test device circuit | 60 |
| Figure 5.6 Optical micrographs of 2 mm CTC cavities in Si: 50%, 60%, 70% PCM fraction (clockwise from top left) | 64 |
| Figure 5.7 Optical Micrograph of Si test chip with film heater and window for IR camera access | 67 |
| Figure 5.8 Schematic of the experimental setup with regenerative cooling and (a) an enclosed jet or (b) an unconfined jet for cooling the TEC hot electrode | 69 |
| Figure 5.9 TEC used for CTC regeneration in experiment | 70 |
| Figure 5.10 Fan cooled regeneration setup | 72 |

| | |
|---|----|
| Figure 5.11 Electric diagram of CTC test device circuit with TEC and IR camera | 74 |
| Figure 5.12 Typical CTC temperature response curve during melt and re-solidify transient | 76 |
| Figure 5.13(a) Temperature histories of $\sim 45 \text{ W/cm}^2$ heat flux hotspots on Pyrex substrate with externally attached CTC | 81 |
| Figure 5.13(b) Temperature histories of $\sim 130 \text{ W/cm}^2$ heat flux hotspots on Pyrex substrate with externally attached CTC | 81 |
| Figure 5.14 Comparison of thermal response of $\sim 130 \text{ W/cm}^2$ heat flux hotspots on Pyrex substrate with CTC attached using Kapton tape TIM vs. spreadable (Ceramique) TIM | 83 |
| Figure 5.15 Temperature histories of $\sim 395 \text{ W/cm}^2$ heat flux hotspots with 2 mm diameter CTCs monolithically integrated as a part of the device under test, both compared to a Si baseline | 84 |
| Figure 5.16 Comparison of temperature histories of $\sim 395 \text{ W/cm}^2$ heat flux hotspots with 2 mm and 3 mm 50% PCM fraction CTCs monolithically integrated as a part of the device under test at 395 W/cm^2 | 85 |
| Figure 5.17 Temperature histories of $\sim 300 \text{ W/cm}^2$ heat flux hotspots with CTCs monolithically integrated as a part of the device under test, compared to a Si baseline | 87 |
| Figure 5.18 Duty cycle of test chips with CTC and attached TEC for regeneration (a) full cycle (b) close up of regeneration portion of cycle | 90 |
| Figure 5.19 Impact of TEC activation time on overall duty cycle | 92 |
| Figure 5.20 Temperature history of device with contact resistance at the TEC-CTC interface | 93 |
| Figure 5.21 Temperature history of fan regenerated CTC | 94 |

| | |
|--|-----|
| Figure 5.22 Temperature history of liquid jet regenerated CTC | 95 |
| Figure 6.1 Schematic of computational domain for CTC model validation (a) 50% PCM fraction, (b) 70% PCM fraction | 97 |
| Figure 6.2 Schematic of a prototype CTC showing the relevant dimensions used in the numerical simulation | 99 |
| Figure 6.3 Comparison of model vs experiment for 50% PCM fraction CTC | 100 |
| Figure 6.4 Comparison of model vs experiment for 70% PCM fraction CTC | 101 |
| Figure A.1 Process flow diagram for Pyrex test chip fabrication | 115 |
| Figure A.2 Process flow of the major stand-alone CTC fabrication steps | 119 |
| Figure A.3 Process flow of major fabrication steps for monolithically integrated CTC and heater/RTD test chip | 127 |

NOMENCLATURE

Abbreviations

| | |
|--------|---|
| BC | Boundary Condition |
| Bi | Biot number |
| BLT | Bond line thickness |
| C | Celcius |
| CTC | Composite Thermal Capacitor |
| CTE | Coefficient of Thermal Expansion |
| Cu | Copper |
| CVD | Chemical vapor deposition |
| D | Diameter of computational domain |
| DAC | Data acquisition system |
| DUT | Device under test |
| E-beam | Electron beam |
| HMDS | Hexamethyldisilazane |
| ICP | Inductively coupled plasma |
| IR | Infrared |
| ITRS | International technology roadmap for semiconductors |
| LPCVD | Low pressure chemical vapor deposition |
| nm | Nanometer |
| Nu | Nusselt number |
| PCM | Phase change material |

| | |
|---------------|--|
| Pr | Prandtl number |
| PTFE | polytetrafluoroethylene |
| PT | Plasma Therm |
| R | Electrical resistance |
| \mathcal{R} | Radius |
| RCA | Radio Corporation of America: clean process named after the company where it was first developed |
| Re | Reynolds number |
| RF | Radio frequency |
| RIE | Reactive ion etch |
| RPM | Rotations per minute |
| RTD | Resistance temperature detector |
| St | Stefan number |
| SSC | Solid state cooler |
| T | Temperature |
| TEC | Thermoelectric cooler |
| TIM | Thermal interface material |
| V | Volt |
| W | W |
| cm | Centimeter |
| c_p | Specific heat |
| h | Heat transfer coefficient |
| h | Enthalpy |

| | |
|----------|---------------------------------------|
| h_{sl} | Latent heat of solidification |
| mA | Milliamp |
| mJ | Millijoule |
| ms | Millisecond |
| mT | Millitorr |
| q'' | Heat flux |
| sccm | Standard cubic centimeters per minute |
| t | Time |
| z | Thickness |

Greek symbols

| | |
|---------------|--|
| α | Under relaxation factor |
| β | Local liquid fraction |
| ε | Emissivity |
| κ | Thermal conductivity |
| μm | Micrometer |
| ρ | Density |
| Σ | Effective cross sectional area for calculating spreading |
| τ | Time constant |
| θ | Non-dimensional temperature |
| Ω | Ohm |

Superscripts

| | |
|----|---------------|
| ° | Degrees |
| qs | Quasi- steady |

Subscripts

| | |
|----------|--|
| ∞ | Far field ambient conditions |
| cs | Cross sectional area |
| eff | Effective value |
| hs | Hotspot |
| melt | Refers to the melting temperature |
| n | Current iteration number (numerical simulations) |
| ref | Reference value |
| sense | Refers to the sensing device (e.g. the sensing resistor used to determine current) |
| Si | Silicon |
| test | Refers to the test device (Pt heaters) |
| throttle | Conditions at the point where the device has reached its throttle criterion (e.g. maximum allowable temperature) |

SUMMARY

While 3D stacked multi-processor technology offers the potential for significant computing advantages, these architectures also face the significant challenge of small, localized hotspots with very large heat fluxes due to the placement of asymmetric cores, heterogeneous devices and performance driven layouts. In this thesis, a new thermal management solution is introduced that seeks to maximize the performance of microprocessors with dynamically managed power profiles. To mitigate the non-uniformities in chip temperature profiles resulting from the dynamic power maps, solid-liquid phase change materials (PCMs) with an embedded heat spreader network are strategically positioned near localized hotspots, resulting in a large increase in the local thermal capacitance in these problematic areas.

Theoretical analysis shows that the increase in local thermal capacitance results in an almost twenty-fold increase in the time that a thermally constrained core can operate before a power gating or core migration event is required. Coupled to the PCMs are solid state coolers (SSCs) that serve as a means for fast regeneration of the PCMs during the cool down periods associated with throttling events. Using this combined PCM/SSC approach allows for devices that operate with the desirable combination of low throttling frequency and large overall core duty cycles, thus maximizing computational throughput. The impact of the thermophysical properties of the PCM on the device operating characteristics has been investigated from first principles in order to better inform the PCM selection or design process.

Complementary to the theoretical characterization of the proposed thermal solution, a prototype device called a “Composite Thermal Capacitor (CTC)” that monolithically integrates micro heaters, PCMs and a spreader matrix into a Si test chip was fabricated and tested to validate the efficacy of the concept. A prototype CTC was shown to increase allowable device operating times by over 7X and address heat fluxes of up to $\sim 395 \text{ W/cm}^2$. Various methods for regenerating the CTC have been investigated, including air, liquid, and solid state cooling, and operational duty cycles of over 60% have been demonstrated.

CHAPTER 1: INTRODUCTION

1.1 Introduction

The tremendous rate of growth in performance capability of electronics devices over the past few decades has been accompanied by the introduction of some significant thermal challenges including power consumption, heat generation and large non-uniformities in chip temperature profiles [1, 2]. Localized hot spots with heat fluxes exceeding $200\text{-}300\text{ W/cm}^2$ have become more common with chip architectures that cluster high power units on the processor to minimize overall chip size [3]. This has typically required thermal management systems that must be designed not only to handle the large background heat fluxes, but also localized hotspots.

Much of the current work involving on-chip hotspot cooling utilizes solid-state refrigeration using thermoelectric coolers (TEC), for example, [4-7]. An alternative solid state cooling (SSC) approach that does not use traditional bulk thermoelectric elements is the use of thin film thermoelectric elements or superlattice coolers [8, 9]. These thin film SSCs offer some advantages in terms of heat fluxes dissipated and improved integration within traditional electronics. However, SSC technology still cannot currently dissipate the largest hot spot heat fluxes, which approach 1 kW/cm^2 .

An alternative to SSC cooling that holds promise is direct liquid cooling. Liquid cooling is the most energy efficient of the chip cooling approaches available [10]. Furthermore large heat fluxes can be addressed with liquid cooling. Recently, localized heat fluxes in excess of 500 W/cm^2 were removed using evaporative cooling [11]. Heat fluxes in excess of 350 W/cm^2 were dissipated using a liquid jet within a compact system

that eliminated the need for an external radiator- a key challenge of implementing liquid cooling [2]. Many of the methods currently being considered for applying direct liquid cooling to hotspots are reviewed in [12].

While liquid cooling is efficient and effective, a key limitation to implementing it as a cooling choice in electronics systems is the packaging concerns. Liquid based systems must prevent leakage, the need for an external regeneration of the coolant, and sometimes require large and costly pumps. Addressing these thermal challenges will be further complicated in next generation 3D architectures that rely on stacking multiple microprocessors or electronics devices on top of one another to achieve enhancements in computing performance. Architectures that vertically integrate the cores in a 3D multi-tier package allow for a number of additional design advantages, including shorter wire lengths, increased packaging density, and heterogeneous technology integration that translate into a range of potential performance benefits such as decreases in noise, capacitance, and power consumption [13]. In these 3D architectures, lack of access to the internal tiers of the 3D stack will make integrating hotspot cooling solutions such as those investigated in [2, 12, 14] increasingly difficult.

As an alternative to active hotspot cooling, computational control schemes such as Dynamic Core Migration (DCM) can potentially levelize the thermal profiles across the chips by actively migrating computations from hotter to cooler areas of the die to keep any one area of the chip from overheating [15]. DCM schemes are enabled by the move towards multi core and many core microprocessors in response to Pollack's and Amdahl's scaling laws. Pollack's and Amdahl's scaling laws indicate that for power-constrained chip designs, architectures that implement many simple, low power cores

should maximize the system's overall performance-per-W, as long as the code is massively parallelizable [16].

To avoid limitations in computation speed due to the serial portions of the code, asymmetric core architectures can be implemented where a few higher power serial cores augment the performance of the low power cores to provide additional throughput [17].

While a DCM approach can mitigate some hotspots, nonhomogeneous architectures which introduce dedicated components such as serial cores [15], may still experience hotspots due to their potentially higher heat fluxes, larger size, and decreased redundancy, [18]. To compensate for the higher power densities the serial cores will either experience more throttling events during an intra-migration time slice or higher migration frequencies [19, 20]. A recent study of different multiplexing techniques for reducing maximum chip temperatures showed that very small (sub-ms) inter migration time slices had to be implemented in order to avoid the presence of hotspots [21].

In DCM schemes, there is parasitic computational cost associated with each throttling event that can become significant over time when the cycling is too rapid [22]. In addition to the computational cost, there is a power consumption associated with core migration that increases with increasing migration frequency [23]. Furthermore, rapid thermal cycling can lead to reduced lifetime reliability for the chip [24]. To minimize the performance losses associated with these gating and throttling events, an optimized system should be designed that can operate for longer periods without requiring an idle for cool-down, and have as short of an idle time as possible.

1.2 Proposed Approach and Relevant Work

In order to address the unique challenges associated with thermal non-uniformities in 3D many-core architectures, an approach is proposed that is a departure from the traditional approach of bringing a specialized liquid cooling device to the hotspot to locally enhance heat transfer. Instead of attempting to increase the heat transfer coefficients in the hard-to-access internal layers of a 3D stack, the design proposed in this thesis seeks to locally increase the thermal capacitance in thermally troublesome areas of the chip to maximize the time that a core or device can operate before reaching its thermal threshold.

As shown schematically in Figure 1.1(a), for dynamically operated micro-architectures, increasing the local thermal capacitance of a device can significantly decrease the required frequency of core hopping, gating, or throttling events. This in turn reduces the parasitic computational overhead associated with the DCM implementation. Thus, matching a device's thermal capacitance to its intrinsic dynamics of power dissipation can "homogenize" the thermal time scales of devices with very different power dissipation profiles.

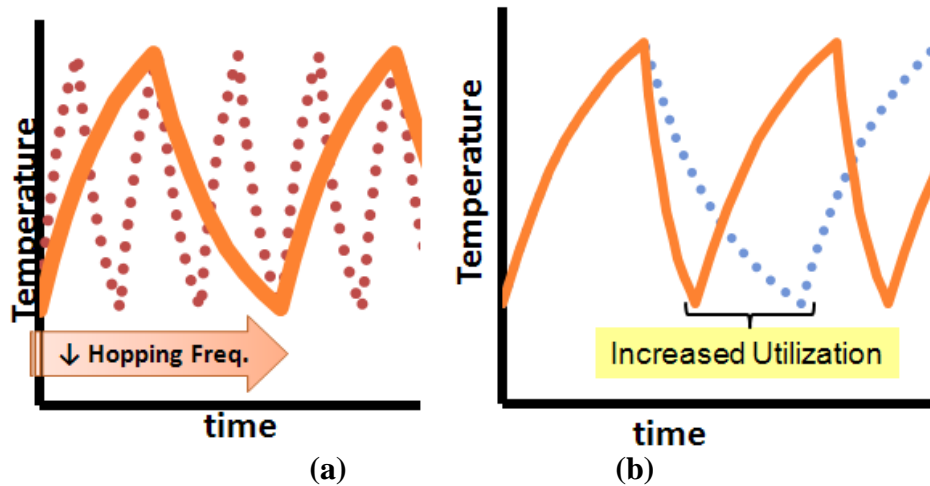


Figure 1.1(a) Impact of increased thermal capacitance on core hopping frequency (b) Impact of regenerative cooling on device utilization

In order to locally alter the dynamic thermal response of the devices, a portion of the silicon on the inactive back side of the chips can be etched away and a material with a higher thermal capacitance, for example solid- liquid phase change materials (PCMs), can be placed in the cavity created by removal of silicon (Figure 1.2). The PCMs, named because of their ability to reversibly melt/solidify during heating/cooling processes, can absorb a large amount of thermal energy at a relatively constant temperature. One challenge of utilizing certain PCMs is that their typically low thermal conductivities limit the amount of material that can be melted prior to the device reaching its threshold temperature. This can be mitigated by using a “composite thermal capacitor” (CTC), consisting of PCM incorporated into a high thermal conductivity matrix to enhance heat spreading and therefore improve PCM utilization.




Figure 1.2 Schematic of CTC integration in a 3D chip stack

While the addition of phase change materials may extend the operating times of dynamically operated devices, some portion of each duty cycle where the device is idled or throttled will inevitably need to be allocated to allow the PCM to cool and re-solidify before beginning the next cycle. Because the CTC stores a significant amount of energy during the melting process, a system with CTC enhancements alone may have relatively long intrinsic cool-down times. Recognizing this potential drawback, a practically viable thermal solution must be designed with a means to rapidly regenerate the CTC to give the computational device the desirable combination of extended operating times and high duty cycles, where duty cycle is defined as:

$$\%duty\ cycle = \frac{device\ operating\ time}{operating\ time + idle\ time}.$$

Solid-state coolers (SSC's), such as thin film thermoelectric or superlattice coolers, are attractive candidates for PCM regeneration, as they are capable of moving very large heat fluxes, as long as minimal sub-cooling is required [9]. In 3D architectures, an added benefit of SSC's comes from the ability to establish physical separation between the cooled and heated junctions through careful placement of the anode and

cathode [8]. Thus the cooled anode can be located within the 3D stack near the CTC to assist in the regeneration of the CTC, while the heated cathode can be placed closer to the external heat sink where it can more readily reject the generated heat to the ambient. Because of the nature of the melting and solidification process, little sub-cooling is required to regenerate the CTC, so that the SSC essentially acts as a sort of thermal fast lane, increasing the rate at which the thermal energy from the CTC travels to the heat sink over spreading and conduction through the stack alone.

While there is certainly application space for significant research into using SSCs as standalone hotspot cooling devices, for example [8, 9, 25], the standalone SSC's heat dissipation capacity is highly dependent on the hotspot size[9], and it must have a cold junction in very close physical proximity to the heat source to be effective [25]. The CTC relaxes those constraints, because it is the CTC that is in close contact with the hotspot directly absorbing the heat load, and the SSC acts only as a secondary thermal transport enhancer with no direct effect on the maximum temperatures the device experiences. Furthermore, because the SSC only operates during the regeneration portion of the operating cycle, the power consumption of the SSC is less of a concern than if it were operating continually as a primary cooling device.

Both the CTC and SSC can be manufactured using standard batch microfabrication techniques, making the proposed solution amenable to the level of high volume integration that is needed for devices in the consumer electronics market.

While SSCs are attractive candidates for CTC regeneration in 3D, traditional cooling methods such as air or liquid cooling remain potential candidates for CTC regeneration, depending on the application and packaging requirements. Devices that

have large heat fluxes and low duty cycle requirements, for example high power RF field effect transistors, may be able to operate with less stringent duty cycle requirements that could be satisfied with an air cooled solution [26, 27]. Furthermore, in 2D applications, bringing a liquid cooled regeneration mechanism into contact with the CTC may be feasible to maximize system efficiency.

There has been significant research in the past into the use of solid-liquid phase change energy storage as an electronic thermal management technique. Many of the approaches have been for applications at the package or heat sink level. For example paraffins were investigated for enhancement of aircraft radar systems over two decades ago [28]. More recently, metal alloys have been investigated as potential PCMs for thermal management applications in ref. [29, 30]. In each case, the PCMs were used at the chip scale, addressing time scales on the order of minute, not milliseconds.

Embedding the PCM within the die to address the rapid time scales and large heat fluxes associated with individual on chip devices is a new approach pioneered through this investigation. An alternative approach for embedding PCMs into the substrates of electronics devices was recently introduced by Tang et al for the management of GaN power amplifiers [31]. While Tang et al considered embedding metallic PCM into the die, no composite structure or regeneration scheme was considered, resulting in relatively low achievable duty cycles.

The design and optimization of composite materials for enhanced heat conduction is a field that has had much interest. Organic PCM composites containing carbon nanotubes [32] and silver nanowires [33] have been considered recently with organic PCM resulting in an approximately 30% enhancement in thermal conductivity. Even with

this enhancement however, the effective conductivity of the composites remains much smaller than a metallic PCM without enhancement. Expanded foams are another common means of enhancing thermal conduction in PCMs with graphite being a good candidate to form the foam makeup [34]. Ultimately, while paraffins have some benefits including being electrical insulators and large latent heats (per unit mass), their intrinsic thermal conductivities may be too small to get to a level where extremely large heat fluxes can be addressed, even with a well-designed composite.

While nanoparticles and foams are typically used in a random dispersion of high thermal conductivity particles in a PCM, methods do exist for identifying an ordered distribution of materials for composite structures. One means of accomplishing ordered or structured composites is through constructal optimization, first introduced by Bejan [35]. The constructal theory attempts to copy the highly ordered patterns found in natural distribution networks such as rivers, tree root and leaves and systematically reconstruct them in systems designed to enhance flow of heat, liquids, or any other analogous substance. The literature on constructal optimization for heat transfer is vast, and well-reviewed in [35], however very little of the literature focuses on transient or phase change applications. One study of a constructally optimized system for solid to liquid phase change was performed by Wang et al [36]; however it focused on a system that addressed much lower heat fluxes and time scales (time constants of $\sim O[10s]$ to $\sim O[1000s]$) than would be useful for this application.

CHAPTER 2: SPREADING IMPACT ANALYSIS

A key feature of the CTC design philosophy is the enhancement of lateral spreading and energy storage into the PCM by improving its effective thermophysical properties, specifically the thermal conductivity (κ) of the composite matrix. Before proceeding with an analysis of how to most effectively maximize the effective thermal conductivity (κ_{eff}) of the CTC, it is valuable to first examine how κ_{eff} , along with the other relevant PCM properties- density (ρ), specific heat (c_p), and latent heat of solid to liquid phase change (h_{sl})- affect the physics of the problem and, in turn, the achievable device operating times. Concentrating on the contribution of lateral spreading to the achievable enhancements in device operating times will allow a more informed decision on whether design of the overall CTC should focus just on the area directly above the hotspot, or on using a larger cross-sectional area accessible through thermal spreading.

2.1 Model Description

A simple model that can be used to study the impact of lateral spreading on device operating times is an annular region of PCM surrounding a cylindrical block of Si of radius \mathcal{R}_{si} , and height z , as shown in Figure 2.1. At the bottom of the Si region is a heat flux boundary condition that represents a localized hotspot. Because the PCM is confined to the annular region at the periphery of the hotspot, this arrangement highlights what can be gained from lateral spreading specifically.




Figure 2.1 Model Schematic for Cylindrical Spreading Analysis

If the domain is thermally thin such that gradients in the vertical direction can be neglected, it can be modeled as 1-D and in the Si region the heat equation can be written in cylindrical coordinates as:

$$\frac{\kappa_{Si}}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T}{\partial r} \right) + \dot{E}_g = \rho_{Si} c_{Si} \frac{\partial T}{\partial t} \quad 0 < r < \mathcal{R}_{Si} \quad t > 0 \quad (2.1)$$

$$\kappa_{Si} \frac{\partial T}{\partial r} = 0 \quad (r = 0, t) \quad (2.2)$$

$$T(\mathcal{R}_{Si}, t) = T_{\mathcal{R}_{Si}}(t) \quad (2.3)$$

Here \dot{E}_g represents the source terms in the Si region; e.g. heat added from the hotspot:

$$\dot{E}_{g,Si} = \frac{q''_{hs}}{z} \quad (2.4)$$

$T_{\mathcal{R}_{Si}}(t)$ is the temperature at \mathcal{R}_{Si} , the boundary between the Si and PCM regions.

From inspection one can deduce the limiting scenarios for PCM performance with respect to the problem formulation described above. In the case of a poorly designed PCM, no energy can be stored in the PCM and the boundary condition equation (2.3) would be changed to an adiabatic condition- $\frac{\partial T}{\partial r}_{r=\mathcal{R}_{Si}} = 0$. In the best case scenario, the PCM would require no wall superheat above T_{melt} in order to drive the melting process. In this case, the boundary condition equation (2.3) would be changed to a constant temperature boundary condition- $T(\mathcal{R}_{Si}, t) = T_{melt}$.

In the PCM region the heat equation is written:

$$\frac{\kappa_{PCM}}{r} \frac{\partial}{\partial r} \left(r \frac{\partial T}{\partial r} \right) = \rho_{PCM} c_{p,PCM} \frac{\partial T}{\partial t} \quad \mathcal{R}_{Si} < r < \mathcal{R}(t), \quad t > 0 \quad (2.5)$$

$$T(r, 0) = T_{melt} \quad (2.6)$$

$$-\kappa_{PCM} \frac{\partial T}{\partial r}_{r=\mathcal{R}_{Si}} = q_{PCM}'' \quad (2.7)$$

$$T(\mathcal{R}(t), t) = T_{melt} \quad (2.8)$$

The location of the melt front is determined from the coupled problem:

$$\rho_{PCM} h_{si} \frac{\partial}{\partial t} (\mathcal{R}(t)) = -\kappa_{PCM} \frac{\partial T}{\partial r}_{r=\mathcal{R}(t)^-} \quad (2.9)$$

$$\mathcal{R}(t = 0) = \mathcal{R}_{Si} \quad (2.10)$$

Here q_{PCM}'' is the heat flux at the Si-PCM interface, which drives the melting process and $\mathcal{R}(t)$ is the location of the melt front.

The melting problem described by equations (2.5-2.10) above fall into a class of problems called Stefan problems which involve heat transfer or diffusion with a moving boundary [37]. The moving boundary in this problem- the melt front propagating radially

away from the Si-PCM interface, introduces a non-linearity into equation 2.5 that precludes the development of an analytical solution.

In order to simplify the problem and allow solution, one may observe that in many problems involving melting, the rate of melt front propagation is much slower than the diffusion time scale. In these instances the problem is dominated by this melt front propagation time scale and as a result a quasi-steady approximation can be made in the single phase region behind the melt front [38].

The quasi-steady approximation involves dropping the transient term from equation (2.5), which after solving results in a temperature distribution in the melt, and melt front position of [38]:

$$T(r, t) = T_{melt} - \frac{q_{PCM}'' \mathcal{R}_{Si}}{\kappa_{PCM}} \ln\left(\frac{r}{\mathcal{R}(t)}\right) \quad \mathcal{R}_{Si} \leq r < \mathcal{R}(t) \quad (2.11)$$

$$\mathcal{R}(t) = \sqrt{R_{Si}^2 + \frac{2R_{Si}q_{PCM}''}{\rho_{PCM}h_{sl}}} \quad (2.12)$$

Substituting the expression for $\mathcal{R}(t)$ into equation 2.10 yields:

$$T_{\mathcal{R}_{Si}} \sim T_{melt} + \frac{q_{PCM}'' \mathcal{R}_{Si}}{2\kappa_{PCM}} \ln\left(\frac{2q_{PCM}'' t}{\mathcal{R}_{Si}\rho_{PCM}h_{sl}} + 1\right) \quad (2.13)$$

The quasi stationary approximation neglects the unsteady term in equation (2.5), and thus neglects any energy storage due to sensible heating. For the approximation to be valid, this contribution should be small relative to the energy storage due to phase change. The Stefan number $St = \frac{h_{sl}}{c_p \Delta T}$ where ΔT is the wall superheat, e.g. $T_{\mathcal{R}_{Si}} -$

T_{melt} , compares the relative contributions of latent and sensible heating in the problem.

When St is large, the quasi stationary approximation is valid.

From equation 2.13 the time scale for melting is:

$$\tau_{melt} \sim \frac{\mathcal{R}_{Si} \rho_{PCM} h_{sl}}{q_{PCM}''} \quad (2.14)$$

while the diffusion time scale is

$$\tau_{diffusion} \sim \frac{\mathcal{R}_{Si}^2}{\left(\kappa/\rho c_p\right)_{PCM}} \quad (2.15)$$

Considering the melt front propagation rate argument that initially motivated the approximation, in order for the melt front propagation rate to be much slower than the diffusion rate, $\tau_{melt} \gg \tau_{diffusion}$. Comparing the two time scales, $\tau_{melt}/\tau_{diffusion} = \kappa_{PCM} h_{sl} / c_p q_{PCM}'' \mathcal{R}_{Si}$, again recovers a large St requirement where $q_{PCM}'' \mathcal{R}_{Si} / \kappa_{PCM}$ represents the temperature difference.

An assumption analogous to the quasi-steady approximation based on time scale comparison is invoked in traditional conduction problems with convective boundary conditions. Following an analysis initially clarified by Bejan, when the convection time scale $\tau_{convection} \sim \rho c_p \mathcal{R} / h$ is compared to $\tau_{diffusion}$ the familiar Biot number, $Bi = h \mathcal{R} / \kappa$ is recovered [39]. For small Bi , the simplified lumped analysis can be used, except in the early regime (small values of $Fo = t/\tau_{diffusion}$) when the thermal wave is propagating across the material [39, 40].

A solution to equation 2.1 for heat generation within a disk with a temperature boundary condition is available [41], however a simplified solution is desirable in order to elucidate the qualitative impact of spreading into the melt on the Si temperatures. Excepting the early regime, when $\tau_{diffusion,Si} \ll \tau_{melt}$ the time dependence of the

temperature in the solid at any point within its interior is the same as the time dependence of the wall temperature which is defined by equation 2.13 [42]. The temperature distribution within the Si can be estimated from the steady state solution to equation 2.1 with the temperature boundary condition given by equation 2.13:

$$T_{Si}(r, t) \sim \frac{\dot{E}_g^{qs}}{4\kappa_{Si}} (\mathcal{R}_{Si}^2 - r^2) + T_{\mathcal{R}_{Si}} \quad (2.16)$$

Here, \dot{E}_g^{qs} , which determines the slope of the temperature increase from the Si-PCM interface ($T_{\mathcal{R}_{Si}}$) to the temperature at the hotspot, is:

$$\dot{E}_g^{qs} = \frac{2q_{PCM}''}{\mathcal{R}_{Si}} \quad (2.17)$$

The implication of this relationship is that because $\tau_{diffusion,Si} \ll \tau_{melt}$, the temperature distribution within the solid can be estimated by a quasi-steady state temperature distribution governed by the rate of energy storage within the Si.

In order to close the relationships above, an estimate must be obtained for q_{PCM}'' . An energy balance on the Si yields:

$$\rho_{Si} c_{Si} \frac{\partial T}{\partial t} = \frac{q_{hs}''}{z} - \frac{2q_{PCM}''}{\mathcal{R}_{Si}} \quad (2.18)$$

Then $\partial T / \partial t$ can be estimated from $\partial T_{\mathcal{R}_{Si}} / \partial t$ (holding q_{PCM}'' constant) to yield:

$$\rho_{Si} c_{Si} \frac{q_{PCM}''^2 \mathcal{R}_{Si}}{\kappa_{PCM} (2q_{PCM}'' t + \mathcal{R}_{Si} \rho_{PCM} h_{sl})} + \frac{2q_{PCM}''}{\mathcal{R}_{Si}} - \frac{q_{hs}''}{z} = 0 \quad (2.19)$$

which is a quadratic equation in q_{PCM}'' that can be solved directly to yield:

$$q_{PCM}'' = \frac{\mathcal{R}_{Si} \kappa_{PCM} (q_{hs}'' t - \rho_{PCM} h_{sl} z)}{z (\rho_{Si} c_{Si} \mathcal{R}_{Si}^2 + 4 \kappa_{PCM} t)} + \dots$$

$$\frac{\sqrt{\mathcal{R}_{Si}^2 \kappa_{PCM} (\rho_{Si} c_{Si} \mathcal{R}_{Si}^2 \rho_{PCM} h_{sl} z q_{hs}'' + \kappa_{PCM} (q_{hs}'' t + \rho_{PCM} h_{sl} z)^2)}}{z (\rho_{Si} c_{Si} \mathcal{R}_{Si}^2 + 4 \kappa_{PCM} t)} \quad (2.20)$$

2.2 Numerical Modeling

While the simplified analysis described above is useful for highlighting some of the important relationships governing the impact of lateral spreading on the device operating times, it only provides an approximate representation of the overall temperature histories. Furthermore, there are key assumptions, such as the large St requirement, that must be satisfied for the analysis to be valid. To provide a point of comparison for the results of the simplified analysis, as well as to study the behavior of materials that fall outside of its range of validity, a more detailed numerical model has been developed to study the problem.

A cylindrical computational domain, shown in Figure 2.1, has been constructed to study the influence of the relevant thermophysical properties of the PCM composite on electronic device performance. The computational domain is axisymmetric about the z -axis and the Si block has dimensions $R_{Si} = 500 \mu m$, $z = 150 \mu m$, and $\mathcal{D} = 1.5 mm$. It is heated from below by a $q_{hs}'' = 500 W/cm^2$ heat flux. Surrounding the Si block is a donut-shaped region of PCM considered to have a range of possible values of effective ρ , c_p , h_{sl} , and κ shown in Table 2.1.

Table 2.1 Thermophysical properties evaluated for cylindrical spreading investigation

| Property | Value Range |
|----------|------------------------------------|
| κ | 1-100 [W/m-°C] |
| c_p | 1 [kJ/kg-°C] |
| h_{sl} | 1-500 [kJ/kg] |
| ρ | 100-10,000 [kg/m ³] |

Except for the hotspot, the remaining external boundaries of the system are considered adiabatic, representing for example, a device embedded in a 3D stack with a large thermal resistance between the device and the global heat sink. This complete thermal isolation condition (the worst case scenario) will be relaxed later in the analysis, but it serves as an appropriate starting point to best illuminate the fundamentals of the local thermal capacitance enhancement process before the added complexity of the packaging and integration related thermal inputs are included in a more comprehensive analysis.

The entire domain is initially set at the solid liquid-transition temperature T_{melt} when the 500 W/cm² hotspot is activated at t=0. The non-dimensional temperature history (θ) of the hotspot is then tracked according to $\theta = \frac{T-T_{melt}}{T_{throttle}-T_{melt}}$, where $T_{throttle}$ is the maximum allowable junction temperature. When $T_{throttle}$ is reached, the device will be have to be throttled or shut-off to allow it to cool down. The generality of the analysis is achieved by tracking θ instead of the dimensional temperatures, as the specific values of T_{melt} do not affect the dynamics (time constant) of the heating process. Instead,

only the temperature budget ($T_{throttle} - T_{melt}$) - determined by the application- is needed. For this analysis, a temperature budget of 15 °C is assumed.

The impact of the melting of the PCM is accounted for by incorporating a modified enthalpy term in the energy equation, as described in [43].

$$\frac{\partial(\rho H)}{\partial t} = \nabla \cdot (\kappa \nabla T) + \dot{E}_g, \quad (2.21)$$

The modified enthalpy term H , accounts for both sensible and latent heating and is defined as:

$$H = \mathcal{h} + \Delta H \quad (2.22)$$

Here the sensible enthalpy, \mathcal{h} , and latent heat ΔH respectively are given by:

$$\mathcal{h} = \mathcal{h}_{ref} + \int_{T_{ref}}^T c_p dT \quad (2.23)$$

$$\text{and } \Delta H = \beta h_{sl} \quad (2.24)$$

In the above equations, \mathcal{h}_{ref} is the enthalpy at the reference temperature T_{ref} , c_p is the constant pressure specific heat of the material, h_{sl} is the latent heat of solidification and β is the local liquid fraction, defined as:

$$\beta = \frac{T - T_{solidus}}{T_{liquidus} - T_{solidus}} \quad (2.25)$$

Here $T_{solidus}$ is the solidification temperature and $T_{liquidus}$ is the melting temperature. For materials that melt and solidify at a single temperature, the β term is undefined at the melting temperature, so a ‘‘mushy zone’’ approximation is made to allow for solution of equation (2.21). The mushy zone approximation assumes the melting process occurs over a range as opposed to a single temperature, beginning at $T_{solidus}$ and completing at $T_{liquidus}$ and vice versa for the solidification process. For the remainder of this

investigation a small $(T_{liquidus} - T_{solidus}) = 0.1$ °C mushy zone is assumed to aid in the convergence of the numerical simulation.

The energy equation is solved using a first order upwind discretization scheme, with a first order implicit transient formulation. To aid in convergence an under-relaxation parameter of $\alpha_\beta = 0.3$ was used for determining β . α_β is defined as:

$$\beta_{n+1} = \beta_n + \alpha_\beta \Delta\beta \quad (2.26)$$

Here n and $n+1$ represent the current and next iteration, respectively, and $\Delta\beta$ is the predicted change in liquid fraction [43]. In addition, an under-relaxation parameter $\alpha_E = 0.9$ was used solving the energy equation. α_E affects the temperature update through:

$$T_{n+1} = T_n + \alpha_E \Delta T \quad (2.27)$$

2.3 Grid and Residual Convergence Investigation

The numerical simulations were solved on a computational domain with 13,799 quadrilateral cells. Grid independence of the solution was evaluated by increasing the number of cells in the computational domain by a factor of four to 55,196 cells, and comparing the predicted operating times before reaching the throttle temperature for the refined and unrefined mesh cases. For the grid convergence test cases, the PCM was assumed to have the properties listed in Table 2.2. When comparing the refined and unrefined cases, there was < 1% difference in the predicted operating times for the two cases. This indicates that simulations solved on the standard grid size of 13,799 cells are sufficiently grid independent.

Table 2.2 PCM material properties evaluated for grid and residual convergence of the numerical simulation

| Property | Value |
|----------|--------------------------|
| κ | 100 [W/m-°C] |
| c_p | 1 [kJ/kg-°C] |
| h_{sl} | 10 [kJ/kg] |
| ρ | 100 [kg/m ³] |

A criterion of scaled energy equation residuals $< 1 \times 10^{-10}$ on the standard grid size was used to indicate convergence of the numerical solution. A reduction of the convergence criterion to residuals $< 1 \times 10^{-11}$ on a test case using the properties in Table 2.2 resulted in a $< 1\%$ difference in the predicted allowable operating times before reaching the throttle temperature. As a result, the energy equation residuals $< 1 \times 10^{-10}$ criterion was considered to be sufficiently small to ensure convergence.

2.4 Results and Interpretation

Figure 2.2 compares the two limiting cases described in section 2.2 (adiabatic BC for poorly performing PCMs and constant temperature BC for ideal PCMs) to the results of numerical solution. Each of the thermophysical properties considered has some ability to move the performance of the PCM towards one the extremes. For example, if κ is small the PCM behaves as an insulator. Similarly if ρ or h_{sl} are small the PCM has little capacity to store heat, again moving the PCM performance towards the adiabatic limit. At the other extreme, if the density or latent heat is infinitely large, the melt front will never move from the Si/PCM interface, thus the problem reduces to the heating of the Si block

with a constant temperature boundary condition at its periphery. At somewhat smaller values of ρ or h_{sl} , the device can still operate near the constant temperature limit if $\kappa \rightarrow \infty$, because negligible temperature gradient will be required to drive the melt front.

In the scenario considered, an enhancement in operating time through spreading of $\sim 2X$ can be achieved by moving from PCMs that operate near the adiabatic limits to PCM composites that operate near the constant temperature limit. However, this is dependent on the allowable $T_{\text{throttle}} - T_{\text{melt}}$. At larger allowable gradients the achievable gains would be significantly larger as the trajectories of the adiabatic and constant temperature curves diverge further from one another.

The results of the numerical simulations shown in Figure 2.2 compare PCMs with $St = 20/3$ to these limiting cases. PCMs with small effective $\kappa \cdot \rho$ combinations approach the adiabatic boundary condition limit, while at $\kappa = 100 \text{ W/m}\cdot\text{C}$ $\rho = 10,000 \text{ kg/m}^3$, the temperature history is almost identical to that of the constant interface temperature case for the values of Fo considered. Over longer times the large $\kappa \cdot \rho$ case would diverge from the constant interface temperature case as the melt front moved further away from the Si/PCM interface.

The $\kappa \cdot \rho$ relationship mentioned and described qualitatively above arises from equation 2.12. For small values of $\frac{2q_{PCM}''t}{\mathcal{R}_{Si}\rho_{PCM}h_{sl}}$, $\ln\left(\frac{2q_{PCM}''t}{\mathcal{R}_{Si}\rho_{PCM}h_{sl}} + 1\right) \sim \frac{2q_{PCM}''t}{\mathcal{R}_{Si}\rho_{PCM}h_{sl}} + 1$ and the inverse relationship of temperature with $\kappa_{PCM} \cdot \rho_{PCM}$ becomes apparent.

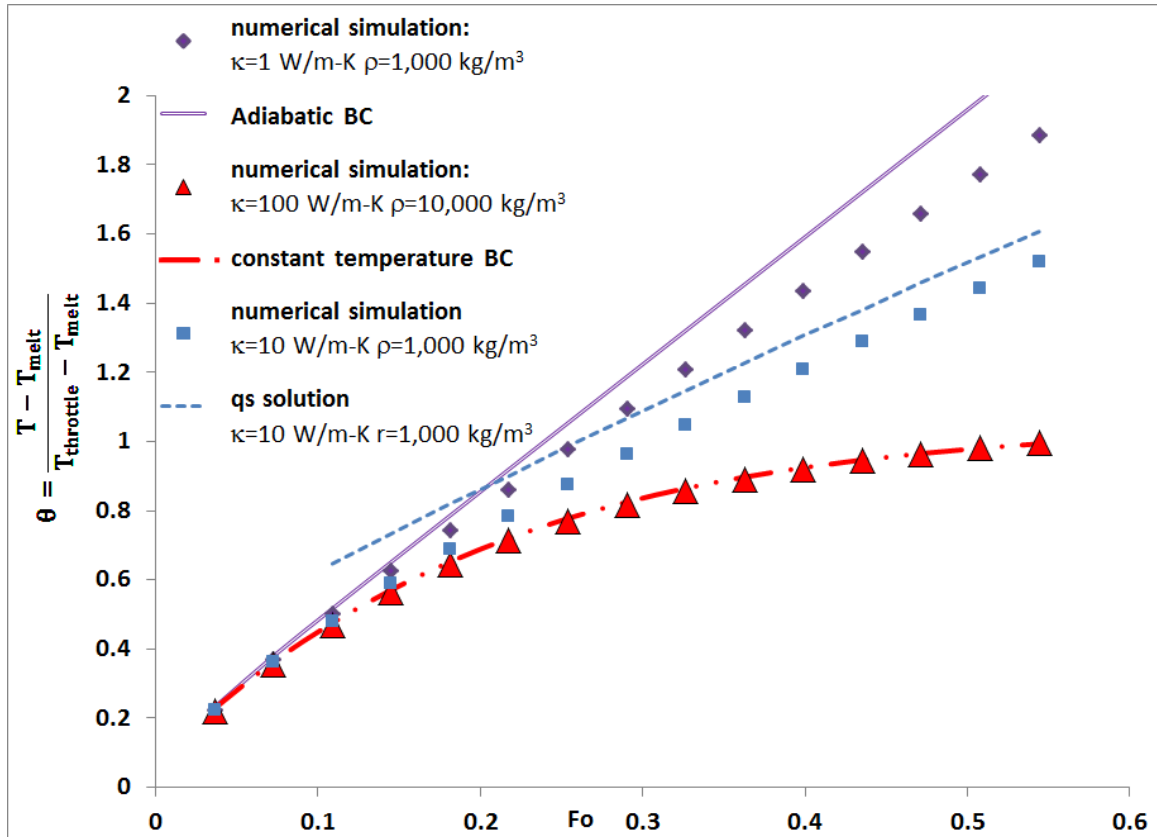


Figure 2.2 Limits on performance modulation of device operating times through spreading into a PCM alone for $h_{sl}=100$ kJ/kg and $c_p=1$ kJ/kg-°C

Also shown in Figure 2.2 is a comparison of the predictions of the quasi-steady analytical model to the results of the numerical solution for $h_{sl}=100$ kJ/kg and $c_p=1$ kJ/kg-°C ($St = 20/3$). At this value of St , the quasi-steady model has reasonable agreement with the numerical solution in the later part of the transient. In the early regime, ($Fo < \sim 0.2$) the solutions converge as the thermal wave propagates across the Si block and the properties of the PCM have little impact. However in the regime where the quasi-steady model is applicable (i.e. large St , and $Fo > \sim 0.2$) the reduced order model gives good insight into how the effective properties of the PCM affect the achievable device operating times.

Figure 2.3 shows the influence of effective κ and ρ on the temperature histories for $h_{sl}=100$ kJ/kg and $c_p=1$ kJ/kg- $^{\circ}$ C . Improvements in device operating times can be achieved through increases in κ , ρ or both. The strength of the forcing for each parameter is similar for both κ and ρ such that configurations with the same value of $\kappa \cdot \rho$ have similar temperature histories, and associated operating times. Increasing κ reduces the magnitude of the thermal gradient needed to drive the melt front a given distance R , while increasing ρ reduces the distance that the melt front travels for a given amount of energy absorbed.

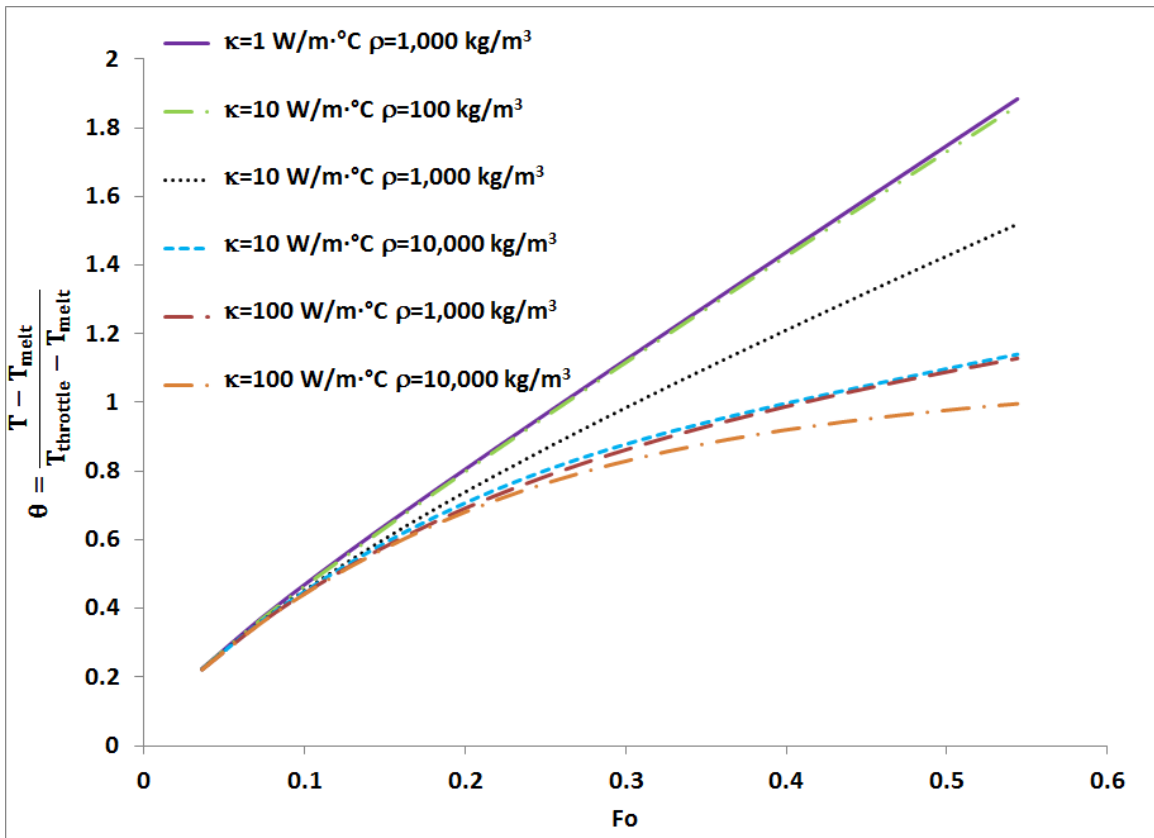


Figure 2.3 Impact of effective thermal conductivity (κ) and density (ρ) on cylindrical spreading temperature histories: $h_{sl}=100$ kJ/kg and $c_p=1$ kJ/kg- $^{\circ}$ C

Considering constant values of effective κ , ρ and c_p ($\kappa=10$ W/m-°C, $\rho=10,000$ kg/m³, $c_p=1$ kJ/kg-°C) Figure 2.4 shows the effect of varying the Stefan number on device performance. The device operating times increase with increasing h_{sl} , although diminishing returns are ultimately observed as h_{sl} becomes sufficiently large. The effect of modulating h_{sl} , is analogous to that of ρ_{eff} - changing the amount of energy stored in a given volume of PCM- although ρ_{eff} has a larger forcing due to its impact on both sensible and latent energy storage.




Figure 2.4 Impact of St on Cylindrical Spreading Temperature Histories for $\kappa=10$ W/m-°C, $\rho=10,000$ kg/m³

While \mathcal{R}_{Si} was not varied parametrically in the numerical analysis, much larger gains in operating times could be achieved with smaller hotspot sizes without increasing

the allowable $T_{\text{throttle}} - T_{\text{melt}}$. The temperature difference between the center of the hotspot where T_{throttle} is observed and the Si/PCM interface scales with \mathcal{R}_{Si}^2 . Because it is the interface temperature that drives the melt front, smaller hotspot sizes would reduce the excess temperature between T_{throttle} and $T_{\text{interface}}$ allowing more energy storage in the PCM before the throttle temperature is reached.

CHAPTER 3: ANALYSIS OF DESIGN SPACE: CTC PERFORMANCE EVALUATION- 2D MODELING

3.1 CTC Layout Description: Si Islands

To evaluate the potential effectiveness of creating a PCM-composite structure to maximize the achievable device operating times, the performance of a CTC has been modeled numerically using two basic layouts. The first layout is comprised of an array of silicon pillars linked thermally with 10 μm wide connectors made of various high thermal conductivity materials including Cu, CVD diamond, and an “ideal” conductor (Figure 3.1(a)). The ideal conductor, whose properties are shown in Table 3.1, is intended to add little thermal mass to the system, while providing a means for near ballistic thermal transport between adjacent Si islands. In practice it could represent, for example, a carbon nanotube (CNT) bundle.


This layout attempts to take advantage of the intrinsically high thermal conductivity of Si, which is already present in the form of the substrate material that would otherwise be etched away to create room for the CTC. The thin 10 μm thermal bridges prevent the Si pillars from being isolated by the low κ PCM that surrounds each pillar.

There are two sub-arrangements considered for the Si-island based layout. Both are shown in Figure 3.1 below. The first arrangement, shown in Figure 3.1(a) has a isotropic 5 x 5 array of Si islands with a center to center spacing of 300 μm . An anisotropic array, shown in Figure 3.1 (b) was also considered where the islands in the 2nd and 4th columns were removed. For each Si-island based arrangement, island sizes of


50 μm x 50 μm , 150 μm x 150 μm , and 250 μm x 250 μm were considered. Each island size is shown in Figure 3.1(b).

Table 3.1 Properties used in simulations for CTC design analysis

| Material | ρ $\left[\frac{kg}{m^3}\right]$ | c_p $\left[\frac{J}{kg \cdot ^\circ C}\right]$ | κ $\left[\frac{W}{m \cdot ^\circ C}\right]$ | h_{sl} $\left[\frac{J}{kg}\right]$ | T_{solid} [$^\circ\text{C}$] | T_{melt} [$^\circ\text{C}$] |
|-----------------|--------------------------------------|--|--|--------------------------------------|-------------------------------------|------------------------------------|
| Si | 2330 | 710 | 150 | | | |
| Cu | 8940 | 386 | 400 | | | |
| CVD diamond | 3500 | 509 | 1800 | | | |
| Ideal conductor | 1 | 10 | 10000 | | | |
| PCM | 9570 | 123 | 40 | 32500 | 60 | 60.1 |



(a)



(b)

Figure 3.1 Schematic of Si-Islands layouts (a) isotropic array (b) anisotropic array- all three island sizes shown (dimensions in m)

3.2 CTC Layout Description: Diamond Shaped Loops


The second layout uses an azimuthal array of diamond shaped loops of high thermal conductivity materials. The diamond shaped loop arrangement for the high thermal conductivity material was first studied by Rocha et al, which showed that this type of layout is effective at enhancing steady state heat conduction in a volume containing two solid materials with differing values of thermal conductivity [44]. The layout was biologically inspired by the venation patterns of leaves in nature.

While a more involved optimization following for example, the constructal theory put forth by [35, 44], might include consideration of varying numbers of loops or irregularly shaped branching structures, this geometry consisting of eight diamond shaped loops was chosen as a design constraint, in order to limit the computational cost associated with the investigation. Eight loops were chosen to align with the inherent symmetry of the problem. That is, the square hotspot can be modeled fully using a 45 degree (1/8) slice of the overall hotspot. An additional benefit of using this geometry as the foundation for the spreader matrix is that experimental prototypes replicating the design can be readily fabricated using standard photolithography and microfabrication techniques.

The size of the regular octagonal computational domain shown in Figure 3.2 is set to fit within a circumscribed 3 mm diameter circle. This limitation on the maximum allowable size of the CTC is set to ensure that it does not monopolize an impractical amount of valuable real estate in the bulk Si, which is also needed for through-layer interconnect routing [45]. The high thermal conductivity materials are laid out in a rhombus shaped region with 45° and 135° internal angles as shown in Figure 3.2. The

rhombus is positioned such that the vertex furthest from the hotspot would be 1 mm from the hotspot center if the region had no thickness (dimension “A” in Figure 3.2). The thickness of the rhombus shaped channels of high thermal conductivity materials is then varied to observe how the trade-off between PCM and high thermal conductivity material volume fraction affects the device operating times.

In each layout, the voids between the spreader materials are filled with a PCM with properties of a hypothetical Bi-In-Sn type alloy with a 60 °C melting temperature (Table 3.1).



Heat-Conductor Spreader Network




Figure 3.2 Schematic of 2D computational domain for CTC modeling- Diamond shaped loops

3.3 Model Description

To simulate the process of activating a high power core and storing the energy in the PCM the modified heat conduction equation (2.21) is solved within a planar (thickness-averaged) CTC computational domain.

The source term \dot{E}_g in equation (2.21) accounts for the heat generation at the hotspot, assumed to be uniformly distributed in the vertical (across the CTC layer) direction. In practice, the heat generation only occurs in a thin layer on the active side of the chip, so this approximation is only valid in the limit of thin PCM layers with high effective thermal conductivity across the plane, such that temperature gradients in the vertical direction are negligible.

The outer boundaries of the computational domain are considered adiabatic, and the heat generation in the hotspot domain is taken as 33 W/mm^3 corresponding to a 500 W/cm^2 hotspot under a $150\mu\text{m}$ thick CTC, in line with projected peak heat fluxes in future devices [20]. For each configuration analyzed, the chips starts at an initially uniform temperature of 59°C (just below the PCM melting temperature) and the simulations are run until the junction temperature reaches 90°C .

Because of the major assumptions included in this model, such as complete isolation from any external heat sink, neglecting temperature gradients in the vertical direction, and large allowable temperature swings (up to 30°C), this initial analysis only provides understanding of what best case potential operating time enhancements can be achieved with the CTC design.

3.4 Grid and Convergence Study

Grid size and residual convergence was tested using several representative cases. In each case tested for grid and residual convergence, diamond was chosen as the high thermal conductivity material. To test the Si-island based arrangements, the grid size was refined and the solution was re-computed for the test cases corresponding to PCM fractions of 58%, 72% and 98%. The 58% and 98% PCM fraction cases correspond to the anisotropic Si- island arrangement (Figure 3.1(b)), while the 72% case corresponds to the isotropic Si-island array. For the diamond shaped loop arrangements, the grid and residual convergence was tested at a PCM fraction of 95%. The standard grid sizes used for the computations are described in Table 3.2. Convergence was tested on the 72% PCM fraction Si island case and the 95% PCM fraction diamond shaped loop case. The scaled energy residuals were found to be converged at a value of 3×10^{-8} for both cases with temperature deviations of <1% with reduction in residual.

Table 3.2 Grid and convergence study for CTC design analysis

| Si Islands | | | | |
|----------------------|----------------------------|---------------------------|-----------------------------------|----------------------------------|
| PCM fraction [%] | Standard Grid Size [cells] | Refined Grid Size [cells] | Operating Time Standard Grid [ms] | Operating Time Refined Grid [ms] |
| 98 | 15,834 | 32,037 | 11.4 | 11.2 |
| 72 | 55,626 | 97,914 | 9.8 | 10.0 |
| 58 | 15,834 | 32,037 | 7.2 | 7.5 |
| Diamond Shaped Loops | | | | |
| 95 | 5,998 | 8185 | 24.5 | 24.4 |

3.5 Results

As shown in Figures 3.3-3.6, there is a significant potential for extending the core operating time before throttling/hopping is required with the use of CTC. When compared with the baseline case of pure Si heat spreading through the device stack, which allows the core to operate for 3 ms before reaching the threshold temperature, an increase in operating time ranging from a factor of 3X for the Si- islands matrix with copper connectors, up to 20X for the diamond-shaped matrix can be achieved.

The Si-islands based arrangements shows a monotonic increase in achievable device operating times with increasing PCM fraction as seen in Figures 3.3 and 3.5. Figure 3.4 shows the location of the melt front when the maximum hotspot temperature reaches 90 °C for the Si-island arrangement with $\kappa=1800 \text{ W/m}\cdot\text{°C}$ connectors. The high thermal conductivity connectors provide an increase in the length of the melt front that increases the amount of PCM that the hotspot can access during the transient. However, the actual Si-islands provide little additional enhancement in the propagation of the melt front. In fact, the melt front never makes it to the next set of Si islands beyond the hotspot before the device reaches its throttle temperature. Much of the area allotted for the CTC in Figure 3.4 remains unmelted when the hotspot reaches 90 °C. While a 3 mm x 3mm cross sectional area was allowed for each Si Island configuration considered during this design analysis, in an actual device the CTC would only be large enough such that majority of the PCM in the CTC is melted during the transient.




Figure 3.3 Time-on-a-core before reaching a 90°C threshold for the isotropic array of Si-islands layout, and different material properties (for quantitative comparison, a baseline of Si only layer yields 3 ms time of operation before reaching a 90°C threshold)

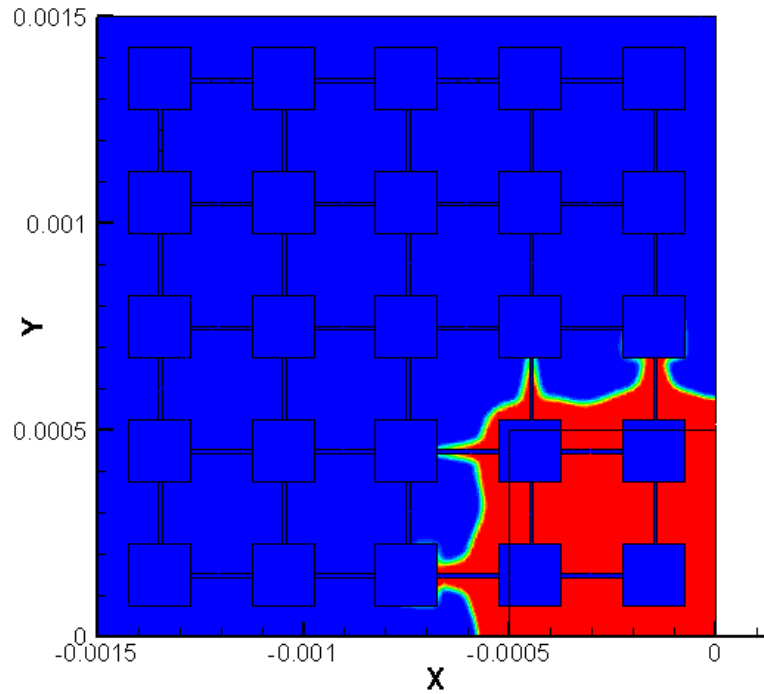


Figure 3.4 Melt front location at throttle time for Si-Islands layout

While the Si islands provide little enhancement in melt front propagation, the isotropic array of Si islands outperform the anisotropic Si-island arrangement. This is because not only were the Si islands removed in the 2nd column of islands, but the high thermal conductivity vertical connectors were removed as well. As was described above, the high thermal conductivity connectors do provide some enhancement in the melt front size, and the loss of those connectors decreases the performance of the overall arrangement.

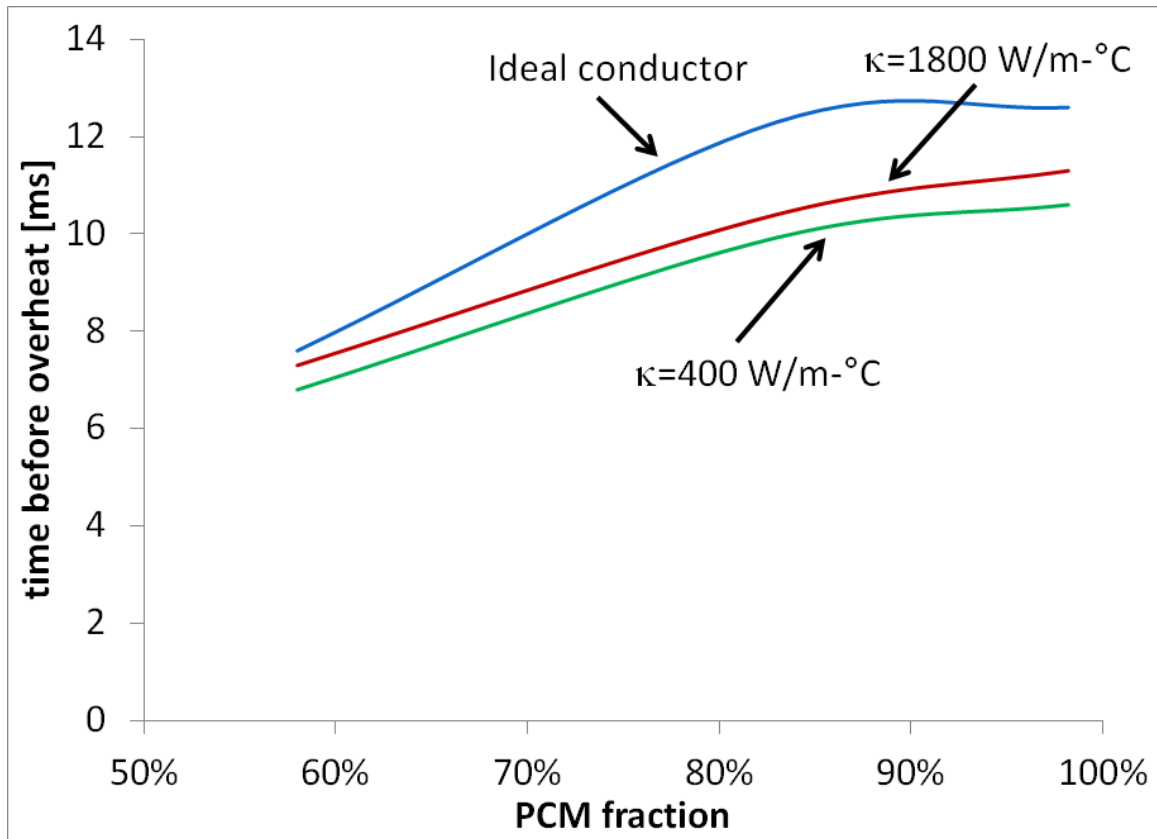


Figure 3.5 Time-on-a-core before reaching a 90°C threshold for the anisotropic array of Si-islands layout and different material properties

When compared with the Si-islands based arrangement (Figures 3.3, 3.5), the diamond shaped loop arrangement consistently delivers longer operating times as seen in Figure 3.6. In the Si-islands arrangement, much of the spreading enhancement comes from the high thermal conductivity materials that connect the individual Si-islands. However, those connectors are only 10 μm wide and make up only a small fraction of the solid material in the CTC. With the diamond shaped loops on the other hand, all the solid material in the CTC is made up of the high thermal conductivity material. The larger fraction of high thermal conductivity material coupled with the larger dimension of the

loops allow the diamond shaped geometry to provide the better overall performance of the layouts considered in this investigation.

Two separate peaks are seen in Figure 3.6. The first local maximum occurs in the regime where the PCM fraction is above 90% and the performance is limited by the tradeoff between the increased spreading that is achieved with adding additional high thermal conductivity materials and the associated decrease in the amount of PCM available for heat storage. With higher κ materials such as diamond, CNTs or graphene, an appreciable increase in spreading can be achieved with the addition of a relatively small amount of the higher κ material, so that the capacitive properties of the composite still closely resemble that of the unaltered PCM. However, with more moderate κ spreading materials such as copper or silicon, the increase in spreading at very high PCM loading is small, diminishing the magnitude of the first peak.

In the second maximum, which occurs at PCM loadings below 60%, the performance is limited primarily by the amount of lateral cross sectional area allotted to the CTC. In this regime, additional spreading could potentially be achieved by adding extra high thermal conductivity materials, but the area that is made available for the placement of the CTC is limited to 3mm in diameter by the electrical design and routing constraints discussed in section 3.2. While the performance of both regimes is governed by the combination of the energy storage due to phase change at the melt front and single phase heating of the composite material in the regions where melting has already occurred, in this lower PCM loading regime the single phase portion of the capacitance is strongly affected by the presence of the higher thermal conductivity materials.

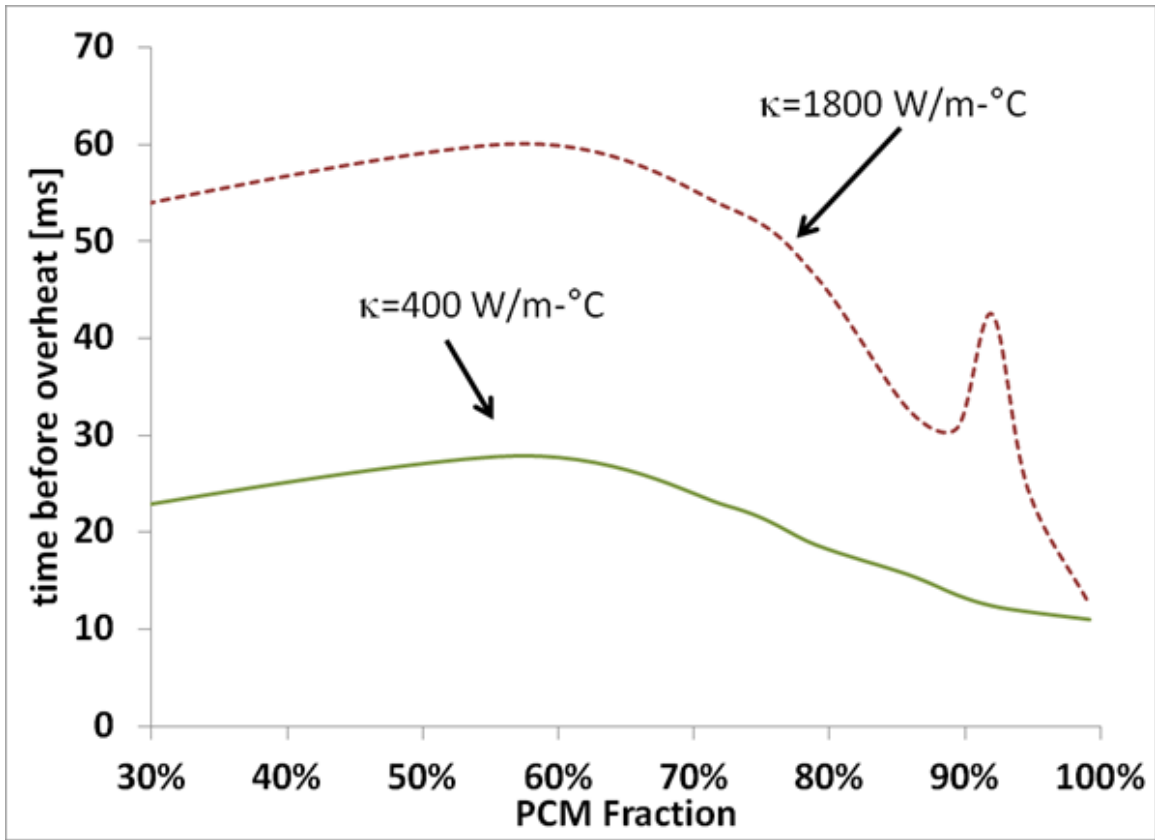


Figure 3.6 Time-on-a-core before reaching a 90°C threshold for the diamond shaped loop layout and different material properties

CHAPTER 4: 3D MODEL-CTC INTEGRATED INTO A 3 LAYER STACK WITH SSC REGENERATION

4.1 Model Description

After considering the potential of various PCM- high thermal conductivity spreader configurations in the 2D domain, a fully three-dimensional model, shown in Figure 4.1, was developed to investigate how the CTC with SSC regeneration would perform when packaged in a configuration similar to what would be experienced in an actual multi-layer stack. The computational domain consists of three 1.5 cm x 1.5 cm x 150 μm silicon tiers with a 20 μm thick polymer between each layer, as a bonding agent. A uniform 100 W/cm^2 background heat flux is generated at the backside of each simulated tier, except in the bottom tier, where in addition to the background heat flux, there is a 1 mm x 1 mm square hotspot in the center.

Directly above the hotspot, a 3 mm x 3 mm x 150 μm portion of the silicon in the bottom tier has been replaced with a CTC with the estimated properties of a Bi-In-Sn alloy with 79 $^{\circ}\text{C}$ melting point (e.g. alloy-174), with an enhanced effective thermal conductivity κ_{eff} ranging from 40 $W/m \cdot ^{\circ}\text{C}$ - 120 $W/m \cdot ^{\circ}\text{C}$ achieved by embedding the high thermal conductivity matrix (e.g., copper or diamond) in the PCM (Table 4.1). The thermal properties of alloy-174 were not available, so the properties of a similar eutectic alloy (*alloy-158*) were used to estimate the properties [46]. Alloys 158 and 174 are similar, with alloy 158 having a slightly lower melting temperature of 70 $^{\circ}\text{C}$. The chemical composition of alloy 158 is $Bi_{.50}Pb_{26.7}Sn_{13.3}Cd_{.1}$, whereas alloy 174 is composed of $Bi_{.75}In_{.26}Sn_{.17}$.




Figure 4.1 Schematic of 3D computational domain and boundary conditions

The expected property differences between the alloys have little importance to this investigation, as the goal is not to concentrate on the performance of a particular PCM, but instead to understand how this class of materials can enhance the

performance of thermally constrained electronics. There are dozens of eutectic and non-eutectic alloys with melting temperatures below the maximum expected temperatures of current generation and ITRS 2024 devices [47]. The ultimate selection of a PCM for the implementation in a device will depend on a number of inputs, including T_{melt} , thermal performance, PCM toxicity (e.g. tolerance of lead or cadmium content), and packaging concerns (e.g. CTE, wettability, materials compatibility) among many other things. Thus the important consideration in modeling the candidate PCM is to ensure that the model properties are representative of materials that may ultimately be implemented.

Above the CTC is a 3 μm thick 3 mm x 3 mm Solid State Cooler (SSC) with the properties of $\text{Si}_{1-x}\text{Ge}_x$ (Table 4.1) which can optionally be activated to enhance the regeneration rate during the cool down portion of the cycle. The cooling potential of thin film SSCs is typically affected by the cross-sectional area of the device, such that thin films SSCs typically have a cross sectional area in the range between $\sim 50 \mu\text{m} \times 50 \mu\text{m}$ and $\sim 100 \mu\text{m} \times 100 \mu\text{m}$ [48]. While this is much smaller than the 3 mm x 3mm size of the SSC modeled here, it is assumed that the 9 mm^2 area of the SSC is populated by a uniform array of 100 $\mu\text{m} \times 100 \mu\text{m}$ individual SSCs, each operating at 31 mA.

Because the SiGe layer must be electrically connected to the substrate above it, and is also in direct contact with CTC below it, the thermal contact resistances between the components in this portion of the thermal path (the CTC to SSC interface, and the SSC to substrate interface) are assumed to be zero. This is in contrast to the interfacial resistance assumed to exist between vertically adjacent surfaces elsewhere in the stack due to the need for an adhesive layer to bond and thermally couple the individual tiers.

This assumption of negligible thermal contact resistance between the CTC-SSC-substrate thermal path is potentially achievable if the SSC is attached to the backside of the tier vertically adjacent to the hotspot as shown in Figure 4.2(a). In this configuration, the SSC is grown directly onto the vertically adjacent tier thus eliminating any contact resistance on that side of the thermal circuit. Where the SSC contacts the CTC, the PCM portion of the CTC (e.g. solder like material) makes up the interface and holds the SSC in place. In other words, the PCM portion of the CTC acts as both a thermal capacitor and a TIM, eliminating any need for an extra interface.

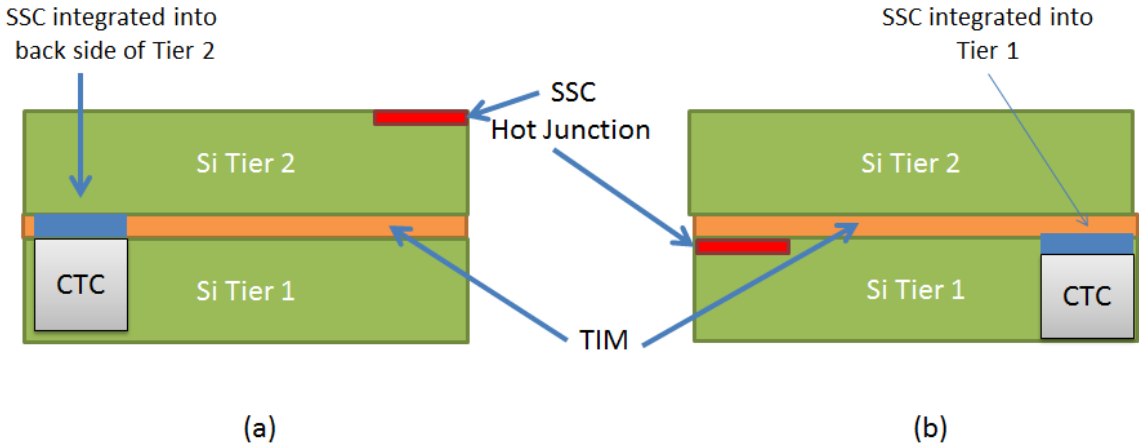


Figure 4.2 SSC integration approaches (a) SSC integrated into vertically adjacent tier (as modeled) (b) SSC is in same tier as CTC

An alternate configuration, shown in Figure 4.2(b), exists where the SSC and CTC are all integrated into a single tier, and then attached to the rest of the stack using a thermally conductive adhesive. In this layout, inter-tier thermal resistance would be approximately uniform, with no areas of zero TIM resistance. This sort of configuration may be simpler from an assembly and packaging perspective, however the integration

of the SSC and CTC into a single tier would remain as a design challenge. Only the configuration of Figure 4.2(a) has been analyzed in this investigation.

The middle Si tier, which acts as the substrate for the SSC, has a vertical Si via at its edge that passes through the top tier. The purpose of this pass through via is to allow the SSC cathode (its “hot” junction) to be cooled at the top of the stack, minimizing the complexity of integrating any liquid cooling solution that may be needed to keep the temperatures within design limitations.

To minimize the numerical cost of modeling this problem, only a single quadrant of the overall chip stack is modeled. Thus a 90° rotational symmetry is assumed. This symmetry condition, as well as the location of the Si via and cathode shown in Figure 4.1, indicate that the current from the SSC is spread evenly between four individual cathodes, each located in one quadrant of the 3D stack.

Also at the top of the stack is a heat sink with 0.05 °C/W thermal resistance for dissipating the global heat load.

In addition to the SSC+CTC geometry described above, a Si-only case is also considered as a baseline point of comparison. The volume occupied by the CTC and SSC shown in Figure 4.1 is replaced with Si, with properties as shown in Table 4.1.

Because the inter-tier thermal interfacial resistance is neglected in this region for the SSC+CTC case, the same assumption is made for the Si-only case. In practice, a Si-only chip stack with no CTC integration would have a uniform inter-tier interfacial resistance, determined by the bonding method. This zero contact resistance assumption was included in the baseline case however, so that the performance impact of the

CTC+SSC could be highlighted independent of any additional benefit that might be obtained because of the differences in packaging approach of the two techniques. Because of this neglected thermal resistance, the performance benefits of the CTC+SSC are considered as conservative predictions.

Table 4.1 Material properties used in 3D computational analysis

| | ρ $\left[\frac{kg}{m^3}\right]$ | κ $\left[\frac{W}{m \cdot ^\circ C}\right]$ | c_p $\left[\frac{J}{kg \cdot ^\circ C}\right]$ | h_{sl} $\left[\frac{kJ}{kg}\right]$ | T_{melt} [$^\circ C$] | S $\left[\frac{\mu V}{^\circ C}\right]$ | σ $[\Omega m]^{-1}$ |
|--------------------------|---|---|---|--|------------------------------|--|-------------------------------|
| Si [49] | 2330 | 150 | 710 | | | 450 | 2.8×10^4 |
| Polymer adhesive [50-52] | 1190 | 10 | 1500 | | | | |
| SSC [7, 53, 54] | 3078 | 8.5 | 602.5 | | | 234 | 5.8×10^4 |
| Anode/Cathode [43] | 8978 | 387.6 | 381 | | | -1.9 | |
| PCM [46] | 9580 | 40-120 | 184 | 39.8 | 79 | | |

4.2 Model Implementation and boundary conditions

Equation 2.21 is employed to model the physics of the heating and cooling process, however several additional source terms and modified boundary conditions have been added to account for the impacts of the SSC, interfacial layers, and the external heat sink. In each 20 μm polymer layer a $5 \times 10^{10} W/m^3$ source term is applied, to simulate the $100 W/cm^2$ background heat flux of the layer.

At the top of the stack convective boundary condition

$$-k \frac{\partial T}{\partial z} = h(T - T_{\infty}) \quad (4.1)$$

is applied with an ambient temperature T_{∞} of 27 °C and an effective heat transfer coefficient, h , of $87 \text{ kW/m}^2 \cdot \text{°C}$, as shown in Figure 4.1. For the 2.25 cm^2 chip area, this convective boundary condition is equivalent to a heat sink with a thermal resistance of 0.05 °C/W attached to the top of the stack.

Along the exterior edges of the stack an adiabatic boundary condition is applied. Along the centerline a zero flux boundary condition is also applied to account for the symmetry in the problem.

$$\frac{\partial T}{\partial x}(x = L_{chip}, y, z, t) = 0 \quad (4.2)$$

$$\frac{\partial T}{\partial y}(x, y = L_{chip}, z, t) = 0 \quad (4.3)$$

$$\frac{\partial T}{\partial x}(x, y = 0, z, t) = 0 \quad (4.4)$$

$$\frac{\partial T}{\partial y}(x = 0, y, z, t) = 0 \quad (4.5)$$

As mentioned previously a prescribed heat flux boundary condition is applied at the bottom ($z=0$) surface. Aside from the $1 \text{ mm} \times 1 \text{ mm}$ hotspot, everywhere along the bottom surface the background flux is applied according to:

$$-k \frac{\partial T}{\partial z}(x > L_{hotspot}, y > L_{hotspot}, z = 0, t) = 100 \frac{\text{W}}{\text{cm}^2} \quad (4.6)$$

The initial temperature distribution in the 3D stack, $T(x, y, z, t = 0)$, is given by the steady state solution to equation 2.21 with an adiabatic heat flux boundary condition at the hotspot:

$$k \frac{\partial T}{\partial z}(x \leq L_{hotspot}, y \leq L_{hotspot}, z = 0, t < 0) = 0 \quad (4.7)$$

At time $t=0$ the hotspot heat flux is changed to:

$$-k \frac{\partial T}{\partial z} (x \leq L_{hotspot}, y \leq L_{hotspot}, z = 0, 0 < t < t_{throttle}) = q''_{hs} \quad (4.8)$$

Here q''_{hs} is either 500 W/cm^2 or 1 kW/cm^2 depending on the severity of the hotspot scenario being considered. q''_{hs} is applied until the throttle time, $t_{throttle}$, which is the time it takes for the peak hotspot temperature to rise by $11 \text{ }^\circ\text{C}$, at which point the boundary condition is changed to adiabatic, to simulate completely shutting the core down during the cool down portion of the cycle.

$$k \frac{\partial T}{\partial z} (x \leq L_{hotspot}, y \leq L_{hotspot}, z = 0, t \geq t_{throttle}) = 0 \quad (4.9)$$

At $t = t_{throttle}$, the simulation is allowed to proceed in one of two ways to allow for regeneration of the CTC. Either the SSC is left idle and the solution is allowed to proceed with the conditions described above, or the SSC is activated to enhance the CTC regeneration rate. For each CTC effective thermal conductivity considered, both cases are solved so that the regeneration rate enhancements from the SSC can be quantified.

To simulate the SSC operation, additional source terms are included into the heat transfer model of the device, which include the Peltier effect at the anode metal-SiGe, SiGe-silicon, and silicon-cathode metal interfaces according to [9]:

$$\dot{E}_g = I(S_i - S_{i+1})T, \quad (4.10)$$

where I is the supply current to the device, S_i is the Seebeck coefficient of the i^{th} material (i.e. metal, SiGe, etc.) with i increasing in the direction of current flow, and T is the temperature of the interface between the i^{th} and $(i+1)^{\text{th}}$ layers. Because these

Peltier effect terms are modeled using volumetric sources and sinks, each surface is given a nominal thickness of 1 μm over which the source is distributed.

Three additional source terms are added to account for Joule heating in the SSC, silicon substrate, and the vertical silicon via. Joule heating is computed from

$$\dot{E}_{g-joule} = I^2 R, \quad (4.11)$$

where the resistance R is the sum of the one-dimensional conduction resistance,

$$R_{1D} = z/\sigma A, \quad (4.12)$$

the electrical spreading resistance,

$$R_{spr} = 8/3\pi^2\sigma r; r = \sqrt{\frac{\Sigma}{\pi}} \quad (4.13)$$

and an Ohmic contact resistance of $R_{contact} = 10^{-6} \Omega\text{cm}^2$ [53]. $R_{contact}$ is only considered at the SSC-substrate interface. In these expressions, z is the thickness of the layer, σ is the electrical conductivity of the material, A is the cross sectional area of the layer, and r and Σ are the effective radius and cross sectional area of the region where spreading occurs, respectively. Each of these Joule heating source terms are distributed uniformly in the associated region in the computational domain shown in Figure 4.1 (e.g. $\dot{E}_{g-joule,SSC}$ is distributed in the 3mm x 3mm x 3 μm blue region marked ‘‘SSC’’).

In addition to the hotspot temperatures, the total PCM liquid fraction in the CTC is also tracked throughout the transient. The CTC is initially completely solidified at $t < 0$, and the PCM liquid fraction increases while the hotspot is active. During the regeneration portion of the cycle the liquid fraction decreases during solidification until it

eventually reaches 0 again. The time when the liquid fraction reaches 0, t_{final} is considered the end of the operational cycle.

4.3 Grid and Residual Independence Study

The 3D stack was modeled using a standard grid size that consists of an unstructured tetrahedral mesh with 221,900 cells. Grid independence was checked by increasing the grid density by about a factor of four to 860,720 cells, and then even more to 2,691,675 cells. A κ_{eff} of $120 \text{ W/m} \cdot \text{°C}$ was assumed and the problem was solved to ensure that both $t_{throttle}$ and t_{final} remained invariant with increasing grid size. A scaled energy conservation equation residual $< 5.0 \times 10^{-10}$ was used as a convergence criterion in each case. As shown in Figure 4.3 there was a deviation of less than 2% in both $t_{throttle}$ and t_{final} when solved on a refined grid when compared to the predicted values on the standard grid size.

As shown in Figure 4.3, on the refined (4X) grid, the change in $t_{throttle}$ and t_{final} when the energy equations residuals decreased from 1.0×10^{-9} to $< 5.0 \times 10^{-10}$ were small (3%). Thus the solution was considered converged on the refined grid when energy equation residuals decreased to $< 5.0 \times 10^{-10}$.

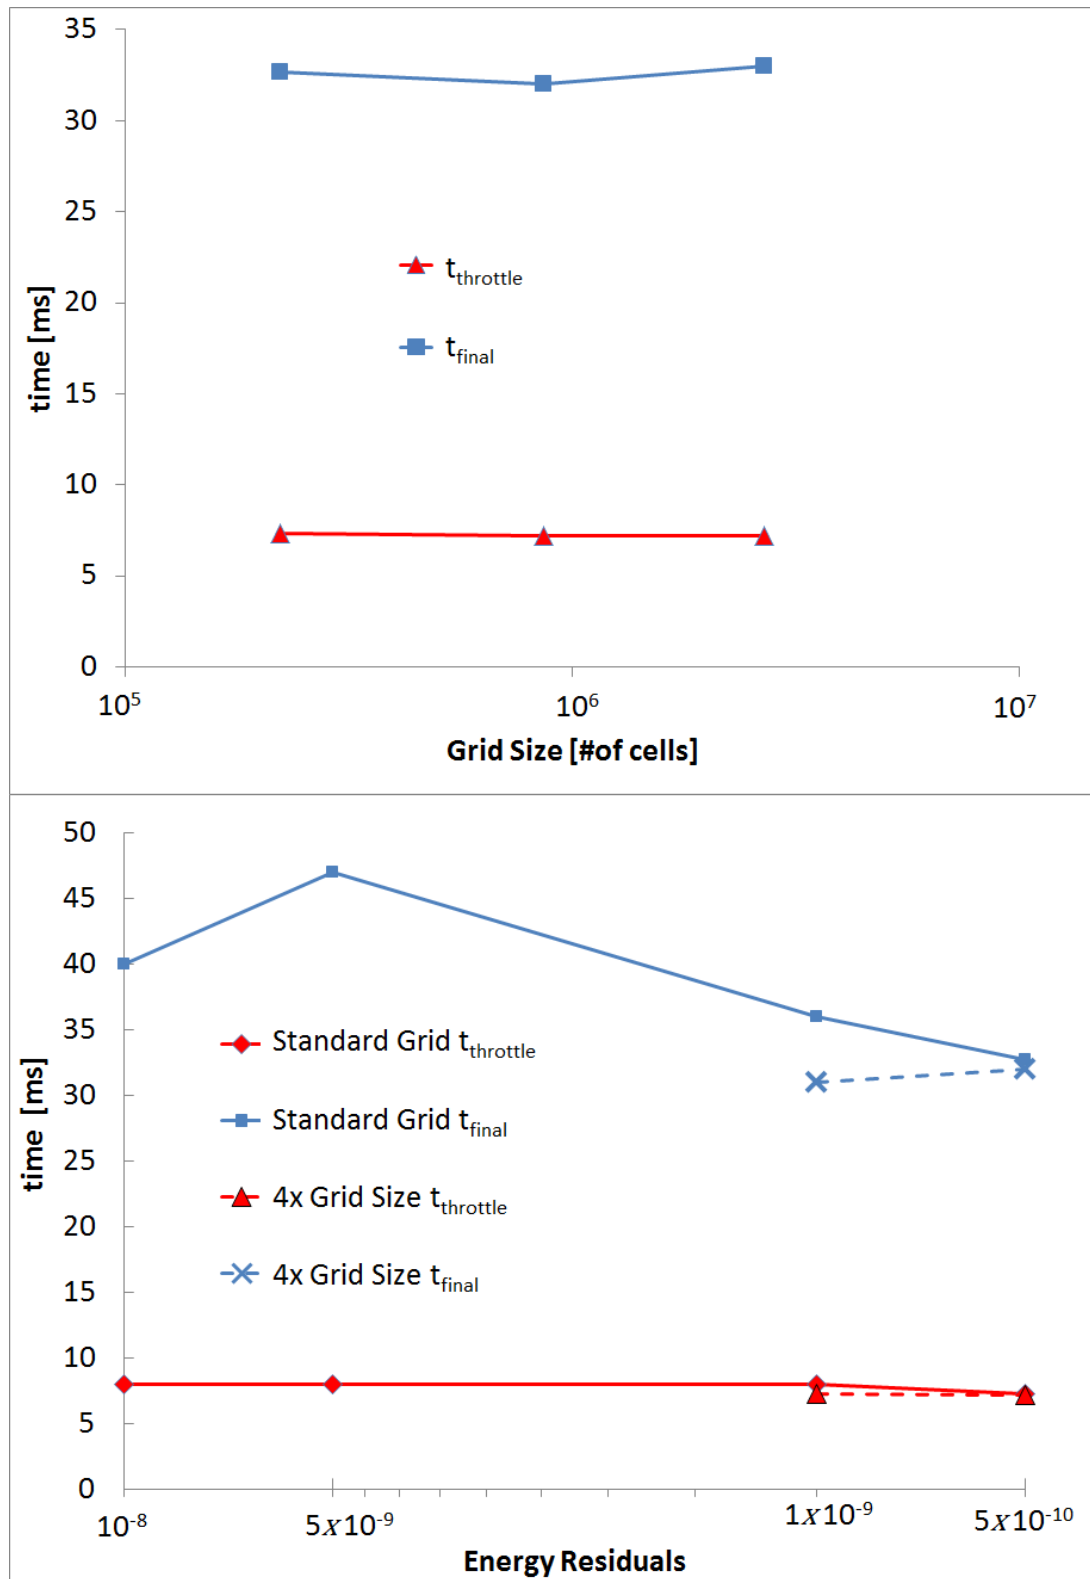



Figure 4.3 Grid and convergence study for 3D computational domain

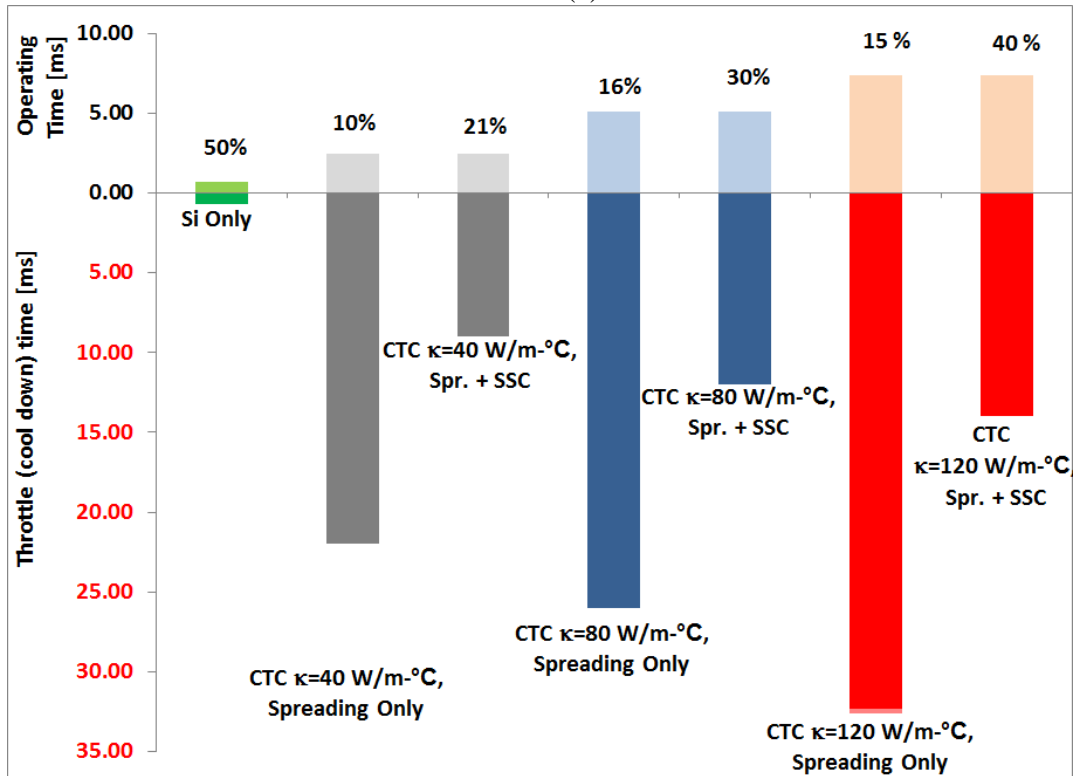
4.4 Results

The results of the 3D simulations (Figure 4.4) show that the potential of the CTC to maximize device operating times when packaged into a realistic device is significant. For the configuration considered, operating time enhancements of over 4X can be achieved at hotspot fluxes of 500 W/cm^2 , and over 10X is possible at 1 kW/cm^2 . Notably, at 1 kW/cm^2 the operating time is increased from 0.7 ms up to 7.5 ms, which would bring the hottest cores in line with the time constant of the rest of the chip [19], significantly simplifying the task scheduling process.

When relying on spreading in the bulk alone for regenerating the PCM, the core's overall availability becomes low, which could be problematic if there is not sufficient core redundancy on the chip. The implementation of active SSC regeneration significantly decreases the unavailability problem, reducing device idle time requirements by 2X- 3X. While the SSC regeneration significantly improves performance of CTC, overall availability of the core with CTC still trails the unaltered Si values for this configuration.



(a)



(b)

Figure 4.4 Results of 3D numerical simulations for CTC performance with and without SSC regeneration: (a) 500 W/cm^2 hotspot heat flux (b) 1 kW/cm^2 hotspot heat flux

An important consideration, however, is that because a device with CTC integration can be designed to match the relevant time constants of the die at large, a potential deficiency in global duty cycle can be addressed through critical core redundancy. A well designed core with CTC integration will be able to operate for at least as long as the scheduled inter-migration time slice. As long as the overall chip is designed such that there are cool cores available to move the thread to at the migration interval, the global duty cycle deficiency can be addressed through this critical core redundancy. On the other hand, if high heat flux device without CTC has its allowable operating time reduced such that it has to be throttled several times within the identified time slice, its duty cycle is capped at its intrinsic duty cycle. The wait times associated with throttling events cannot be addressed through redundancy in this scenario because the thread will already have spent a fraction of the given time slice idled before the scheduler has an opportunity to move the thread to a cooler location. In order to fully take advantage of any designed redundancy, the device must be able to operate for at least the duration of the inter-migration time slice. The levels of redundancy required to account for the intrinsic duty cycles of the CTC are well within the levels being considered in current research. For example, several studies have recently been published considering device redundancies of 25%-75% for next generation many core systems [21, 23].

CHAPTER 5: EXPERIMENTAL CHARACTERIZATION

In order to validate the potential of the CTC concept for enhancing device operating times, several prototype devices have been fabricated and tested. In the first set of experiments, described in section 5.1, a thin film platinum heater was deposited on a Pyrex substrate to simulate a hotspot. A CTC structure that was fabricated on a separate wafer is then coupled to the hotspot through a thermal interface material (TIM) as shown in Figure 5.1(a). This setup is used to observe both the impact of a CTC on device operating times as well as the impact of the TIM on the thermal response of the entire system.

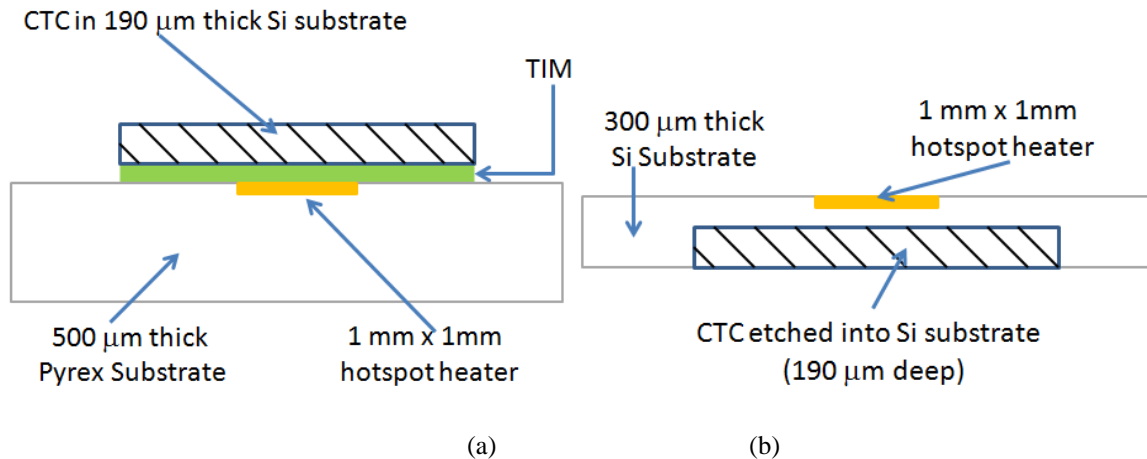


Figure 5.1 Schematics of the experimental setups (a) CTC coupled via TIM to Pyrex test chip (b) Monolithic integration of CTC and test chip in Si

In the next set of experiments, described in section 5.2, the simulated hotspot and CTC are monolithically integrated into a single Si wafer (Figure 5.1(b)). This eliminates the effects of interfacial resistances between the device and CTC and better represents the behavior of a chip that might be integrated into a 3D stack.

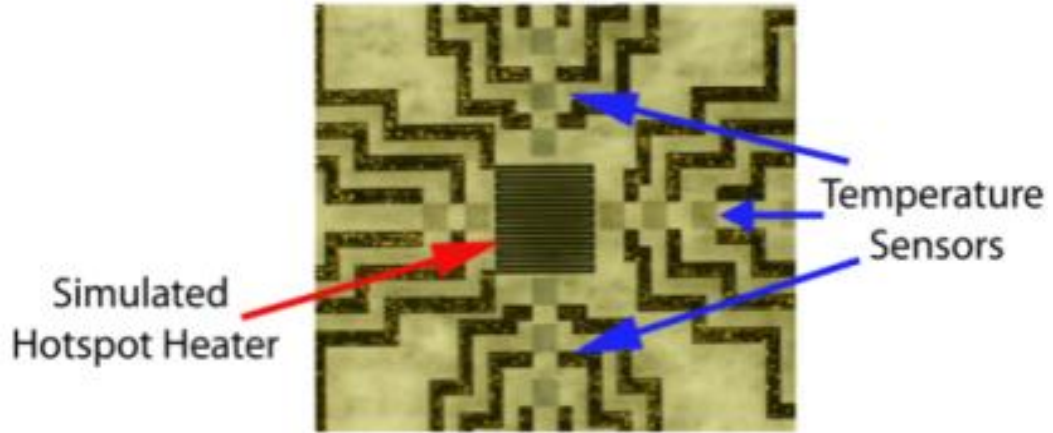
Finally, the impact of the regenerative cooling method on the overall duty cycle is investigated in section 5.3. A commercially available bulk thermoelectric cooler (TEC) was coupled to the CTC shown in Figure 5.1(b) to show how it could potentially enhance the duty cycle of a CTC +Solid state cooler (CTC+SSC) system.

5.1 CTC on Pyrex Test Chip


5.1.1 Device fabrication and experimental setup

A prototype device was fabricated and characterized to validate the potential of the CTC concept for enhancing device operating times. The prototype consisted of two parts: a Pyrex test chip with a 1 mm x 1 mm hotspot heater in the center (a simulated large heat flux device), and a silicon matrix CTC fabricated in a separate Si wafer. In applications where etching the substrate to place the CTC is undesirable, the CTC could potentially be attached to the top (active) side of the device in a separately attached configuration such as this.

The 1mm x 1mm square hotspot heater, and surrounding 250 μm x 250 μm resistive temperature sensors shown in Figure 5.2(a), were fabricated on a Pyrex wafer using a microfabrication process described in appendix A.1. Pyrex was chosen to minimize spreading in the substrate, and to concentrate heating in the vicinity of CTC under investigation. The heater and sensors were made of ~250 nm thick Pt films deposited on a ~25 nm Ti adhesion layer. Pt is chosen because of its linear resistivity vs temperature relationship, allowing the devices to act as not only heaters, but also as resistance temperature detectors (RTDs).



(a)



(b)

Figure 5.2 (a) Hotspot heater layer and RTD layout (b) CTC filled with PCM used in experiments

To calibrate the heaters and RTDs the test chips were heated in an oven and the temperatures and corresponding resistances were measured over a broad range using type T thermocouples and recorded via an Agilent 34970A data acquisition system (DAC). Prior to use, the thermocouples were calibrated in an ice bath, to ensure that they are

accurate within their expected uncertainty range ($\pm 0.5^\circ\text{C}$). Two wire resistance measurements were taken because the resistances of test leads ($\ll 1 \Omega$) were negligibly small when compared to the resistances of the heaters and RTDs ($\sim 1 \text{ k}\Omega$). The resulting temperature vs resistance calibration curves were used to determine the temperature histories of the test devices during operation.

The calibration process showed that resistance of the heaters had a highly linear temperature dependence, as expected. A sample calibration curve is shown in Figure 5.3.




Figure 5.3 Sample Pt heater calibration curve

The prototype CTC shown in Figure 5.2(b) was fabricated by etching diamond shaped cavities in a 190 mm thick Si wafer using an inductively coupled plasma reactive ion etch (ICP-RIE). The cavities were then coated with a 1 μm thick layer of Pt over Ti

adhesion layer to improve wetting of the inner Si cavities by the PCM. The PCM used in the experiments was alloy- 136, a eutectic alloy of $Bi_{0.49}In_{0.21}Pb_{0.18}Sn_{0.12}$ with a 58 °C melting temperature. The major CTC fabrication steps are detailed in appendix A.2.

The CTC shown in Figure 5.2(b) is made of Si with no wetting-promoting coating, and the air gaps and large contact angles near the sidewalls and sharp corners are evident. The Pt coating greatly improved the wettability of the CTC sidewalls, allowing a visibly uniform filling of the cavities. Figure 5.4 shows a CTC partially filled with solidified PCM. The sidewalls and internal angles of the CTC are wetted well by the PCM.

Pt was chosen as the sidewall coating because it typically does not form or tightly adhere to surface oxides which lower the surface energy and reduce wetting. Unlike other noble metals such as Au or Ag, Pt also has an extremely low dissolution rate in most solders which is important to the long term stability of the PCM [55, 56].




Figure 5.4 Partially filled CTC, after Pt coating for improved wettability

5.1.2 Experimental procedure

The CTC was attached to the hotspot using, *C eramique*, a ceramic particle laden spreadable thermal interface material (TIM) as shown schematically in Fig 5.1(a). The entire device was then preheated in a 53  C oven to simulate the elevated temperature conditions that might be experienced inside of a 3D stack. The hotspots were then activated for short pulses of 100 ms at $\sim 45 \text{ W/cm}^2$ and 30 ms at $\sim 130 \text{ W/cm}^2$.

As a point of comparison, the CTC was replaced with an unaltered piece of 190 μm thick Si that was attached to the hotspot in the same manner as the CTC was. This unaltered Si case serves as a baseline to then determine the level of enhancement that is achieved through use of the CTC.

For the pulsed operation the test structures were connected to the test circuit shown in Figure 5.5. The power to the circuit is supplied by an Agilent 6634B DC power supply. A TIP50 NPN type transistor is used as a switch to open and close the circuit downstream of the test device. The base current providing the switching operation of the transistor is supplied by an Agilent 33250A waveform generator, which sends square pulses of the width of the activation time (100 ms or 30 ms) to the device. The edges of the square waves were set at 5 ns ramp time.

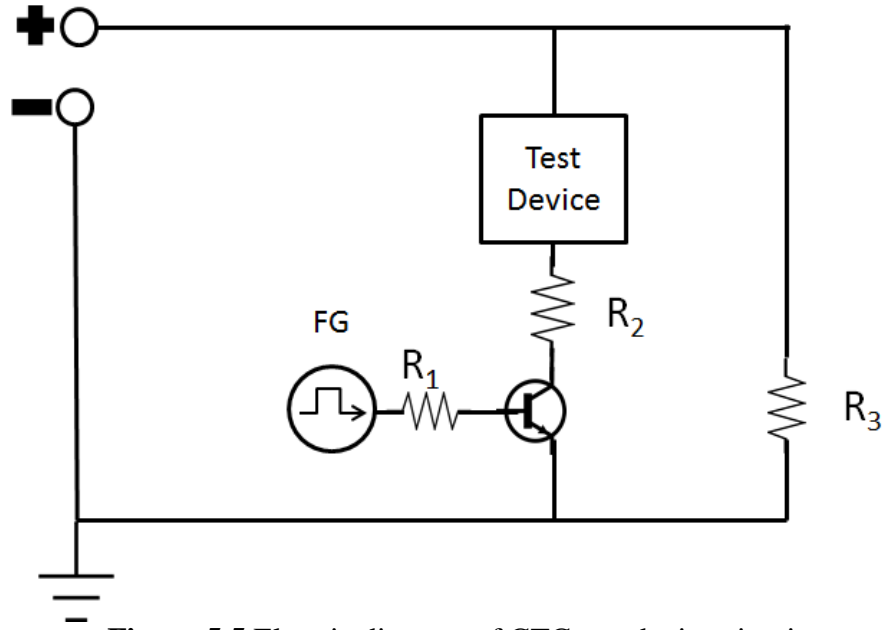


Figure 5.5 Electric diagram of CTC test device circuit

The 82 k Ω bypass resistor, R_3 , shown in the circuit was put in place to allow the power supply to remain at its target voltage with a small amount of current flowing when the test device portion of the circuit is open, thus removing the intrinsic ramp time of the power supply from the temporal response of the main circuit. The sensing resistor, R_2 , was a calibrated 53.5 Ω resistor. The voltage drop across R_2 , V_{sense} , is measured as a function of time using a Tektronix 2014B oscilloscope. The change in current flow through the circuit as the test device heats up can then be determined from V_{sense} and R_2 using Ohm's Law:

$$i_{circuit} = \frac{V_{sense}}{R_2}$$

The small 1.4 Ω resistor R_1 simply provided a small load for the waveform generator to ensure it did not exceed its current rating when the transistor circuit was closed.

With the current and voltage drop across the device known, the resistance of the test device was determined again using Ohm's law, and compared to its calibration curve in order to determine its temperature history throughout the transient. Because of the inherent noise in the measurement circuit, each data point is taken as the average value measured over 16 identical pulses, executed sequentially with sufficient time between each pulse to ensure the device has been fully regenerated before activating the next pulse.

The nominal dissipated heat fluxes for each experiment are determined by computing the Joule heating in the test devices at the beginning of the transient and dividing by the 1 mm x 1 mm cross sectional area of the test heater. The stated heat fluxes correspond to the values present at the beginning of the transient, however, due to the increase in resistance of the heaters with increasing temperature. This effect slowly decreases the current through the circuit and thus the heat dissipated in the test device, although the total change in heat flux over the course of the transient is less than 5% for all the experiments described in this work.

5.1.3 Investigation of TIM impact on device performance

Because the CTC and hotspot heater are coupled through a TIM it is important to consider the impact of that interface on the overall thermal response of the system. Spreadable TIMs typically require one or more atmospheres of pressure to achieve a minimal bond line thickness (BLT) [57]. Because the CTC was fabricated in a thin and fragile (190 μm thick) wafer, the large pressures needed to achieve optimal TIM

performance were not possible to apply to the CTC to minimize the possibility of breaking the device. Instead the TIM was spread as thinly as possible on the Pyrex substrate, and the CTC was then placed on top of the TIM with only light pressure applied manually. While care was taken to spread the particle laden TIM as thin as possible to minimize both its thermal resistance and capacitance, it was not possible to ensure that the BLT was identical from experiment to experiment.

In an attempt to understand and control for the impact of the interfacial resistance on the system response, the particle laden TIM was replaced with a Kapton® based thermal tape to minimize the sample to sample variability in bond line thickness. The thermal tape was comprised of 50.8 μm thick Kapton® tape coated on each side with a 12.7 μm thick layer of Thermaphase® phase change thermal compound. The procedure described above in section 5.1.1 was repeated, and the results were compared.

5.2 Monolithic Integration in Si

While using Pyrex as the substrate for the simulated chip provides a concentrated hotspot and mimics the performance of a device built on a poorly conductive substrate, such as silicon-on-insulator devices, controlling the interfacial resistance between the CTC and Pyrex substrate is very difficult. Because the thermal resistance of this interfacial layer plays such a significant role in determining the system's overall transient thermal response, it must be minimized in order to fully capture the performance potential of the CTC. Given the difficulty of predictably minimizing this interfacial resistance with the externally attached CTC setup described in section 5.1, a monolithic prototype was fabricated with the hotspot and CTC made on the same Si wafer. This not

only helps in highlighting the performance of the CTC, but it also provides a closer representation of the operating conditions that are experienced in a 3D stack.

The monolithic device was fabricated on a 500 μm thick double-sided polished wafer with the simulated hotspot built on one side atop a thin SiO_2 passivation layer. Next, the CTC sidewalls were etched into the back side of the wafer using an ICP-RIE process. The wafer was then thinned to a thickness of 300 μm such that the final device has 190 μm deep CTC on top of a 110 μm thick base. Finally a 1 μm thick layer of Pt was sputtered on the CTC structures to improve its wetting characteristics. The fabrication process is described in detail in appendix A.4

Just as in section 5.1, an unaltered Si device is used for comparison to the CTC performance. For these tests the control device is a 300 μm thick Si test chip with heaters and RTD's on one side and no alterations on the back side.

The CTCs that were fabricated had PCM fractions of 50%, 60% and 70% and a diameter of 2mm. A CTC with 3mm diameter and PCM fraction of 50% was tested as well. Optical micrographs of a representative sample of each device are shown in Figure 5.6.



Figure 5.6 Optical micrographs of 2 mm CTC cavities in Si: 50%, 60%, 70% PCM fraction (clockwise from top left)

After fabrication, the experimental procedure followed was the same as that outlined in section 5.1 although the temperature of the oven was slightly lower at 51 °C. Additionally, because of the elimination of the CTC-hotspot interfacial resistance, experiments could be conducted with higher nominal heat fluxes ranging from ~300 to 395 W/cm² for the monolithic CTC test device.

5.3 Regenerative Cooling

5.3.1 Experimental setup

In addition to extending the achievable operating time of high heat flux devices, rapid regeneration of the CTC is needed to ensure that the devices have large duty cycles—the combination of both long operating times and short idle/regeneration times. In order to demonstrate the regeneration of the CTC using a SSC, several regeneration approaches were tested on the 2 mm 70% PCM fraction device described in section 5.2.

As detailed in chapter 1 and [58], thin film SSC's are attractive candidates for regeneration, especially in 3D or embedded applications where access to the CTC is limited. In addition to SSC, liquid or air cooling are also potential candidates for CTC regeneration for certain applications.

In order to evaluate these regeneration approaches, three different setups were tested. First, a commercially available bulk thermoelectric cooler (TEC) was coupled to the CTC to evaluate the potential of SSC regeneration. Next, air cooling was evaluated using a standard CPU fan for regeneration. Finally, liquid cooling was tested using a 900 μm water jet impingement.

When the simulated hotspot is idle, it is very difficult to accurately determine the temperature of the device using the RTD method, because there is little to no current flowing through the circuit. Instead, the thermal profile is measured using an Infrascopes II infrared (IR) thermal imaging camera. The use of the camera required the test device to be exposed to ambient air, as opposed to the experiments described in sections 5.1 and

5.2, where the device was in a constant temperature oven to simulate a steady-state baseline heating due to a background heat dissipation. In order to preheat the devices to their steady state idle temperatures, a polyimide film heater was attached to the top of the test chip using adhesive as shown in Figure 5.7.

The function of the film heater is to provide background heating and to preheat the substrate before activating the device. The film heater heats the entire device to ~50 °C with the hotspot in the idle state.




Figure 5.7 Optical Micrograph of Si test chip with film heater and window for IR camera access

The IR camera determines the temperature of the device by counting the infrared photons that strike the InSb detector and comparing the measured irradiance with that predicted by the Planck distribution. Infrared radiation is also produced by the ambient

environment, and this can be a significant source of noise in the temperature measurement. One means of improving the signal to noise ratio of the camera is to increase the emissivity of the surface. A surface with a high emissivity produces more photons at given temperature than a surface with lower emissivity. As a result, the ratio of the number of photons at a given temperature originating at the device under test (DUT) to those from the environment will be larger when the DUT has a high emissivity.

The test heaters are made of Pt and thus have an intrinsically low emissivity. In order to minimize the effect of ambient noise on the temperature measurement, the surface of the heater was coated with a thin film of carbon to increase its emissivity (Figure 5.7).

A small 600 μm window is cut into the film heater above the hotspot so that the surface of the heater is exposed to the lens of the IR camera (Figure 5.7). The IR camera is then focused on the surface of the hotspot, and the emissivity, ε , of the surface is determined by comparing the measured IR emission from the surface at a known temperature to that predicted by the Planck distribution. However, the measured emissivity is affected by shadowing from the sidewalls of the film heaters, making determining the accuracy of the measured surface temperatures difficult. This ε dependence was mediated by introducing the nondimensional temperature profile θ , which will be tracked over an operational cycle to determine the dynamics of the device's duty cycle. Here θ is defined as:

$$\theta = \frac{T - T_{initial}}{T_{max} - T_{initial}}$$

The use of the nondimensional temperature eliminates the unnecessary introduction of uncertainty in ϵ without sacrificing the ability to determine the % duty cycle accurately. This is because the change in radiance measured by the camera tracks directly with the change in θ , regardless of the specific value of absolute hotspot temperature.

5.3.1.1 Thermoelectric Cooler Regeneration Setup

A commercially available, eTEC HV56, bulk thermoelectric cooler (TEC) from Nextreme Thermal Solutions was coupled to the CTC described in chapter 5.2, as shown in Figure 5.8 [59]. The dimensions of the cold junction of the TEC, shown in Figure 5.9 are 2.07 mm x 3.08 mm which is slightly larger than the 2 mm diameter circular CTC.




Figure 5.8 Schematic of the experimental setup with regenerative cooling and (a) an enclosed jet or (b) an unconfined jet for cooling the TEC hot electrode

When the TEC was used for CTC regeneration, two different configurations for cooling the hot junction of the TEC were tested. Both hot junction cooling approaches are shown schematically in Figure 5.8. Initially, the hot junction of the TEC was soldered to

the surface of the housing containing a water jet to ensure the TEC received sufficient cooling (Figure 5.8(a)). The cold junction of the TEC was then placed in contact with the CTC with no other interfacial enhancements such as soldering or a TIM. The cold junction of the TEC was not soldered to the CTC because when both surfaces were rigidly attached, even small shear forces would cause the TEC to break.

When the device was tested, it was found that the interfacial resistance between the TEC and CTC in this configuration was large enough that negligible impact on the duty cycle was observed from activating the TEC. Nonetheless, data from this first configuration was still included in this study to show how inefficient cooling can negatively impact overall system duty cycles.




Figure 5.9 TEC used for CTC regeneration in experiment

To minimize the resistance between the TEC cold junction and the CTC that was observed in the experimental setup shown in Figure 5.8(a), the TEC was detached from the water jet and its cold junction was soldered to the CTC using alloy 136 (the same material used to fill the CTC cavities). Because the cold junction was rigidly attached, a

direct impingement of an unconfined water jet was instead used to cool the TEC hot junction (Figure 5.8(b)). The TEC was covered with a piece of 16 μm thick aluminum foil to protect the electric circuit from the water jet. The interface between the aluminum foil and the TEC hot junction in Fig 5.8(b) was filled with the ceramic particle laden TIM described in section 5.1.

The maximum cooling observed using the unconfined jet configuration occurred at 3V. Based on the manufacturer's data [59], it is estimated that this would result in a maximum cooling power of the TEC of $\sim 3\text{W}$ over the 0.064 cm^2 surface area of the cold electrode or a cooling power of about 50 W/cm^2 .

The TEC was operated only when the hotspot heater was off. In order to synchronize the hotspot and TEC operation a second transistor and waveform generator were added to the measurement circuit (Figure 5.11). Because of the potentially higher currents drawn by the TEC, a 2N3055 transistor was used for switching the TEC as opposed to the transistor used for switching the simulated hotspot.

5.3.1.2 Fan Cooling Regeneration

Regeneration using air cooling was investigated using a simple 60 mm x 25 mm 19 CFM fan blown on the backside of the entire chip (Figure 5.10). The chip was separated from the outlet of the fan by a 35 mm spacer. The airflow was allowed to exhaust from the periphery of the chip, which was not sealed. The fan was run continuously throughout the transient, both when the hotspot was powered and unpowered. The Nusselt number based on the fan slot diameter for cooling in an

arrangement such as this can be estimated from [60] to be about 10 in the area beneath the fan hub.




Figure 5.10 Fan cooled regeneration setup

5.3.1.3 Liquid Cooling Regeneration

Regeneration using liquid cooling was tested using a 900 μm water jet that was directly aligned with the hotspot. The water jet was not allowed to directly contact the back side of the chip and CTC however, to ensure that the electrical circuitry and CTC were not damaged. Instead the CTC was covered by a small 4 mm x 4 mm square sheet of 300 μm copper. The copper lid was soldered to the CTC opening using alloy 136. The remainder of the backside of the chip was covered with aluminum foil, which had a 4 mm opening cut in its center that aligned with the copper lid. The gap between the foil and the copper CTC lid was sealed using polyimide film tape.

The liquid jet was operated continuously throughout the transient just as was done in the previously described regeneration setups. When a highly efficient cooling method like liquid cooling is employed, if the size of the localized cooler is not closely matched

to the size of the hotspot, a large portion of the coolant's thermal capacity can be wasted carrying the background heating instead of regenerating the intended CTC [2]. The performance of the water jet can be estimated from [61]:

$$Nu_0 = 1.38Pr^{0.329}Re^{0.489}$$

$$\frac{Nu}{Nu_0} = 0.777 \left(\frac{r_{CTC}}{d} \right)^{-0.5}$$

Here Nu is the average Nusselt number of the jet over its 3mm diameter zone of influence, Nu_0 is the Nusselt number in the jet stagnation zone, r_{CTC} is the radius of the jet zone of influence and d is the jet diameter. For the 3 mm diameter area that is potentially cooled by the 1.6 mL/s liquid jet, this translates to a thermal resistance of only 2.75 °C/W. In order to ensure that the liquid coolant is not inundated with carrying a large amount of heat from the film heaters, it is preheated to 48 °C before being delivered to the chip. Because its temperature is similar to the baseline temperature of the chip, it provides little subcooling to the average chip temperature, better highlighting its impact on CTC regeneration.

5.3.2 Device operation and characterization process

After preheating, the Pt hotspot is pulsed using the same method described in section 5.1.2. The hotspot heat fluxes and the length of the test pulses for each configuration considered are listed in Table 5.1. A second waveform generator (FG₂ in Figure 5.11) was used to synchronize the IR camera with the pulse of the hotspot. Using FG₂, a trigger signal was sent to the camera 10-20 ms prior to activating the hotspot so that the beginning of the transient would be captured by the camera.

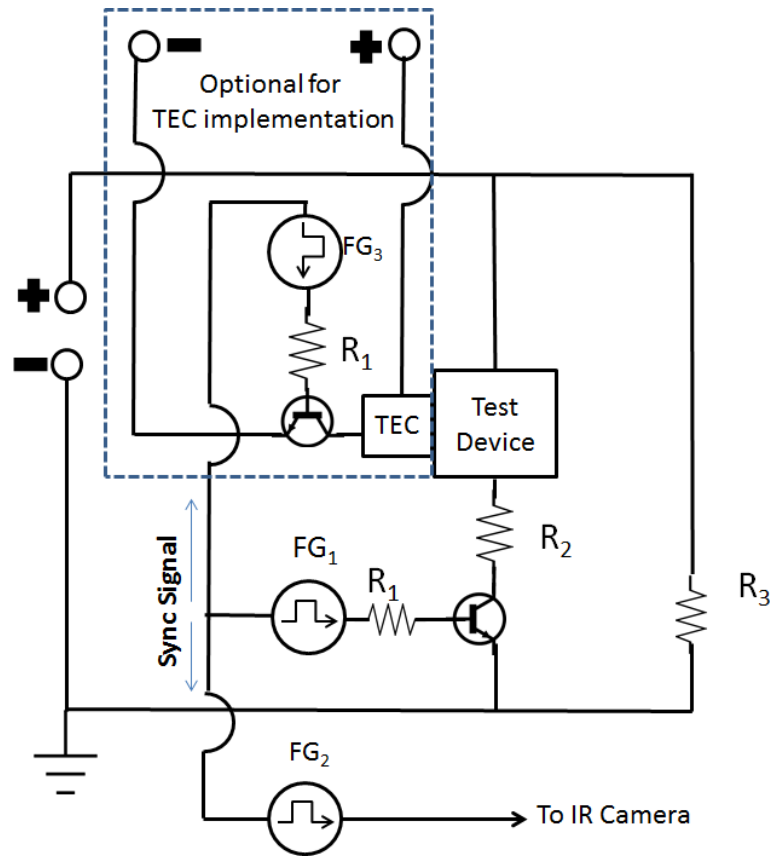


Figure 5.11 Electric diagram of CTC test device circuit with TEC and IR camera

Table 5.1 Pulse length and heat fluxes of the four configurations for CTC regeneration

| Regeneration Approach | Pulse Length [ms] | Heat Flux [W/cm ²] |
|--|-------------------|--------------------------------|
| TEC soldered to CTC, hot electrode cooled with unconfined jet (Figure 5.8(b)) | 250 | 365 |
| TEC not soldered to CTC hot electrode cooled with enclosed jet (Figure 5.8(a)) | 100 | 372 |
| Fan Cooling (figure (5.10)) | 150 | 365 |
| Direct Liquid Cooling (NO TEC) | 250 | 365 |

All of the regeneration approaches considered were operated continuously except for when Thermoelectric Cooler (TEC) regeneration was employed. The TEC was only active when the simulated hotspot was switched off. To synchronize the activation of the

TEC with the deactivation of the hotspot, a third waveform generator (FG₃ in Figure 5.11) was synchronized with the hotspot control circuit. As detailed in Table 5.1, the device cooled with the TEC (TEC soldered to CTC) was pulsed for 250 ms at 365 W/cm². At the 250 ms mark, the device was switched off and either the TEC was switched on, or the stored heat in the CTC was allowed to spread into the substrate for passive regeneration.

To minimize the impact of the ambient noise on the measurements, each data point is taken as the average value measured over 5 identical pulses, executed sequentially with 15 s between each pulse.

A sketch of a typical device temperature history during the melting and solidification process is shown in Figure 5.12. During the regeneration process the device temperature undergoes three typical steps, labeled as such in Figure 5.12. First, immediately after the device is shut off there is a rapid decrease in temperature from the throttle temperature to a temperature slightly above T_{melt} . Next, there is a period of nearly constant temperature solidification, step 2 in Figure 5.12. Finally, when the CTC has re-solidified, the temperature begins to drop again, as the device cools down to its initial temperature. For the purpose of determining the device duty cycles, the test chips are considered regenerated at the last inflection point after subcooling has begun, labeled in Figure 5.12.

The one exception is when TEC regeneration is utilized. If allowed to operate continually, the TEC would subcool the device below its initial temperature. Instead the TEC cooled device is considered regenerated when its temperature has decreased to the level where a non-TEC cooled device would be considered regenerated.

While the device temperature may continue to slowly decrease for several seconds after the point where the device is considered regenerated, the vast majority of the energy will have been rejected from the CTC once the device has re-solidified. In practice, the device could be reactivated after the CTC has been re-solidified when it is mostly regenerated with little difference in overall system performance. For example, for a CTC with a St of 10, the device has rejected $\sim 90\%$ of its stored energy upon re-solidification, and $\sim 99\%$ of its stored energy at the time that its temperature has cooled below T_{melt} to within 5% of its initial temperature.

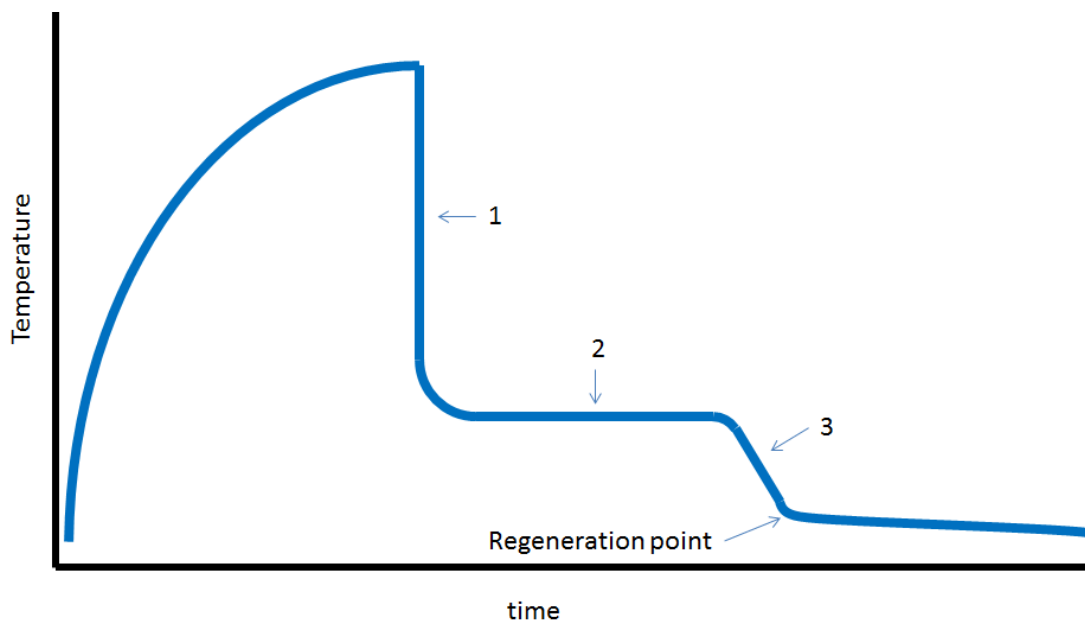


Figure 5.12 Typical CTC temperature response curve during melt and re-solidify transient

5.3.3 Uncertainty Analysis

Measurement uncertainty is determined using an error propagation analysis [62]. The uncertainty of the various components used in the experimental process is shown in Table 5.2. When the RTD method was used for temperature measurement (sections 5.1, 5.2) the measurement uncertainty is dependent upon two uncorrelated uncertainties: the

uncertainty in determining the resistance of the test device, and the uncertainty associated with the temperature vs resistance calibration of the test devices.

The test device resistance, R_{test} , is determined from Ohm's law, as described in section 5.1.2. The determination of R_{test} is dependent upon three measurements: the voltage drop across the test device V_{test} , the voltage drop across the sensing resistor V_{sense} , and the resistance of R_2 (Figure 5.11). V_{test} is determined from the difference between the voltage at the positive terminal of the power supply V_{total} and the upstream terminal of R_2 (see Figure 5.5). With the negative terminal of the power supply grounded, the expression for determining R_{test} can be written as:

$$R_{test} = \frac{V_{total} - V_{test}}{V_{test}/R_2} = R_2 \left(\frac{V_{total}}{V_{test}} - 1 \right) \quad (5.1)$$

V_{total} and V_{test} are uncorrelated, so the uncertainty of their ratio is given by:

$$\frac{\Delta \left(\frac{V_{total}}{V_{test}} - 1 \right)}{\left(\frac{V_{total}}{V_{test}} - 1 \right)} = \sqrt{\left(\frac{\Delta V_{total}}{V_{total}} \right)^2 + \left(\frac{\Delta V_{test}}{V_{test}} \right)^2} \quad (5.2)$$

The total uncertainty can then be determined from:

$$\frac{\Delta R_{test}}{R_{test}} = \sqrt{\left(\frac{\Delta R_2}{R_2} \right)^2 + \left(\frac{\Delta V_{total}}{V_{total}} \right)^2 + \left(\frac{\Delta V_{test}}{V_{test}} \right)^2} \quad (5.3)$$

In equation 5.3 $(\Delta R_2)/R_2$ is the uncertainty of the sensing resistor, which is calibrated by the manufacturer to be accurate with ($\pm 1\%$). $\Delta V_{total}/V_{total}$ is the uncertainty of the power supply ($\sim 0.1\%$) [63], and $\Delta V_{test}/V_{test}$ is the uncertainty of the oscilloscope measurement (3%) [64].

Table 5.2 Uncertainty analysis

| Measurement Tool | Uncertainty [\pm] |
|--|------------------------|
| Oscilloscope | 3% |
| Agilent 34970A data acquisition system | 1 Ω |
| IR camera (noise) | 4% @50 °C 1% @90 °C |
| IR camera (sensor) | 0.1 °C |
| Agilent 6634B DC Power Supply | 50 mV |
| Sensing Resistor [R_2] | 1% |
| Type T thermocouple | 0.5 °C |

The uncertainty of the test device temperature vs resistance calibration is dependent on the uncertainty of the type T thermocouple measurement, the uncertainty of the resistance measurement using the DAC, as well as the error in the linear regression used to estimate the trend of the temperature vs resistance measurements. The uncertainty of the resistance measurements (<0.1%) and the standard error in the slope of calibration regression curve (~ 0.1-0.2% typically) are negligibly small compared to the temperature measurement error (1%), thus the overall accuracy of the calibration curve is $\pm 1\%$.

Substituting into the above expressions, one finds that the measurement uncertainty for the RTD measurements is dominated by the accuracy of the oscilloscope (3%) used for measuring the voltage drop across the sensing resistor (Figure 5.5). As a result, the results of section 5.1 and 5.2 are reported with a 3% accuracy.

When the IR camera is used to determine temperature (section 5.3) the dominant source of uncertainty stems from the ambient IR noise picked up by the detector. As a result, the measurement accuracy increases with increasing heater temperature. At 50°C, the measured fluctuations in IR camera temperature measurement due to noise result in an uncertainty of $\pm 4\%$. At 90°C the measured fluctuations in IR camera temperature measurement due to noise reduce to $\pm 1\%$. There is additional uncertainty in measuring absolute temperature due the need to determine the surface emissivity. However the % duty cycle is determined from the nondimensional temperature history, and is not affected by the emissivity error.

5.4 Results

5.4.1 CTC on Pyrex Test Chip

5.4.1.1 Device Operating Time Enhancements

The results of the experimental characterization of the CTC prototypes, reported in Figure 5.13, show that coupling the hotspot to the CTC enhances the devices' allowable operating times when compared to those observed when the CTC is replaced with unaltered Si. For example, when a 45 W/cm^2 heat flux was applied to the hotspot, increases in operating time before throttling in the range of $\sim 4.5\text{X}-6\text{X}$ can be achieved for allowable temperature increase of $<10 \text{ }^\circ\text{C}$. For higher heat fluxes of $\sim 130 \text{ W/cm}^2$, if an $85 \text{ }^\circ\text{C}$ maximum allowable junction temperature throttle criteria is applied, an increase in operating time of $\sim 4.4\text{X}$ can be achieved.

However, because of the interfacial resistance between the device and CTC, the heat fluxes dissipated are relatively modest when the CTC is not monolithically integrated on a chip. While a 130 W/cm^2 hotspot would be difficult to address with conventional air-cooled approaches [65, 66], in order to address the largest of heat fluxes expected in next generation devices, the TIM resistance must be minimized to allow broad application of two-part CTC approaches. As will be discussed in subsequent sections, eliminating the interfacial resistance through monolithic integration does significantly improve CTC performance, however there remains application space for a stand-alone (via an externally attached) CTC implementation where monolithic integration is not an option.

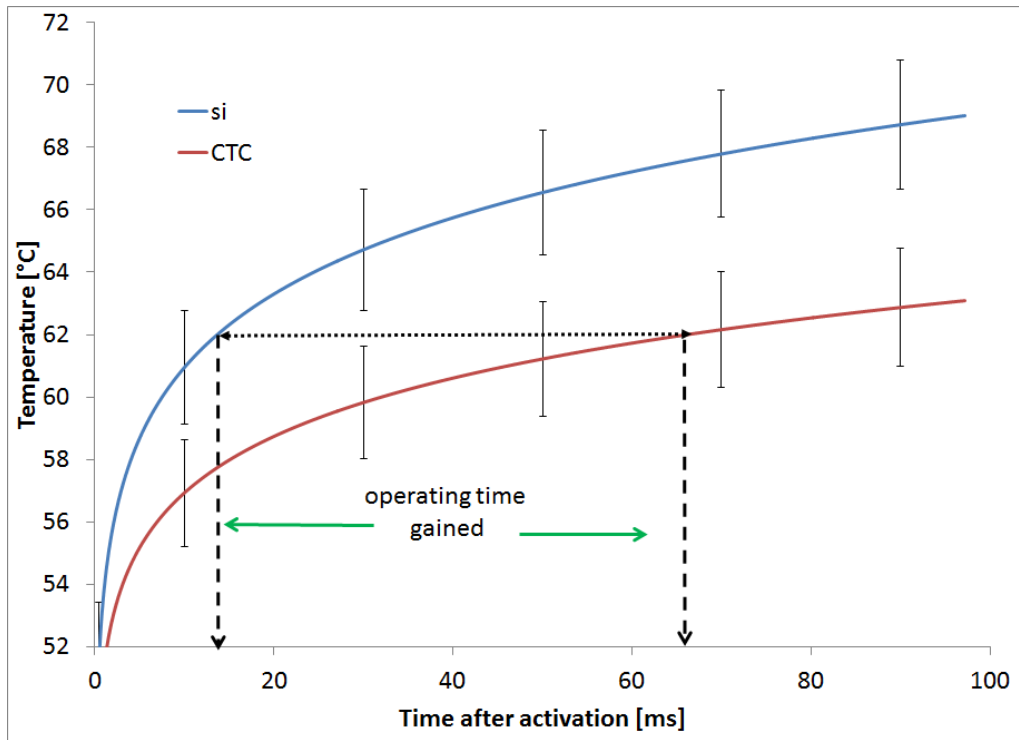


Figure 5.13(a) Temperature histories of $\sim 45 \text{ W/cm}^2$ heat flux hotspots on Pyrex substrate with externally attached CTC

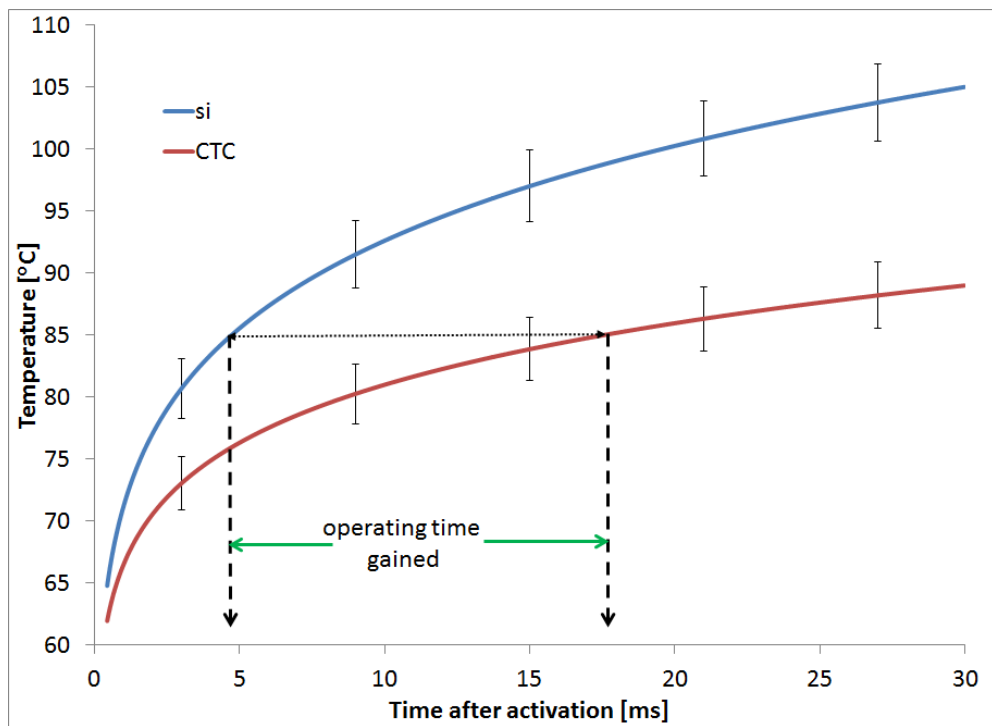


Figure 5.13(b) Temperature histories of $\sim 130 \text{ W/cm}^2$ heat flux hotspots on Pyrex substrate with externally attached CTC

5.4.1.2 Investigation of TIM impact on device performance

When the thermal tape is used the CTC still shows an improvement in device availability over the unaltered Si, as seen in Figure 5.14. However, because the thermal tape has a thermal resistance that is over 10X higher than is achievable with thinly spread particle laden TIM, as well as a higher thermal capacitance, it has a significant impact on the system's overall thermal response. When the thermal tape is in place, the thermal response of the system is dominated by the characteristics of the tape, masking the performance distinctions between the CTC and the unaltered Si. Because this interfacial resistance can be significant, the data from the monolithic CTC prototypes is crucial to fully understanding the potential of the CTC as a thermal management approach.




Figure 5.14 Comparison of thermal response of $\sim 130 \text{ W/cm}^2$ heat flux hotspots on Pyrex substrate with CTC attached using Kapton tape TIM vs. spreadable (Ceramique) TIM

5.4.2 Monolithic Integration in Si

Moving from the separately attached test device to the monolithic design allowed an increase in addressable heat fluxes from 130 W/cm^2 up to $\sim 395 \text{ W/cm}^2$ as shown in Figure 5.15. At $\sim 395 \text{ W/cm}^2$ the 70% PCM fraction CTC provides an enhancement in achievable operating time of the test device of over 650% more than that observed with the unaltered Si test chip. At 50% PCM fraction, the operating time enhancement reduced to 260%.

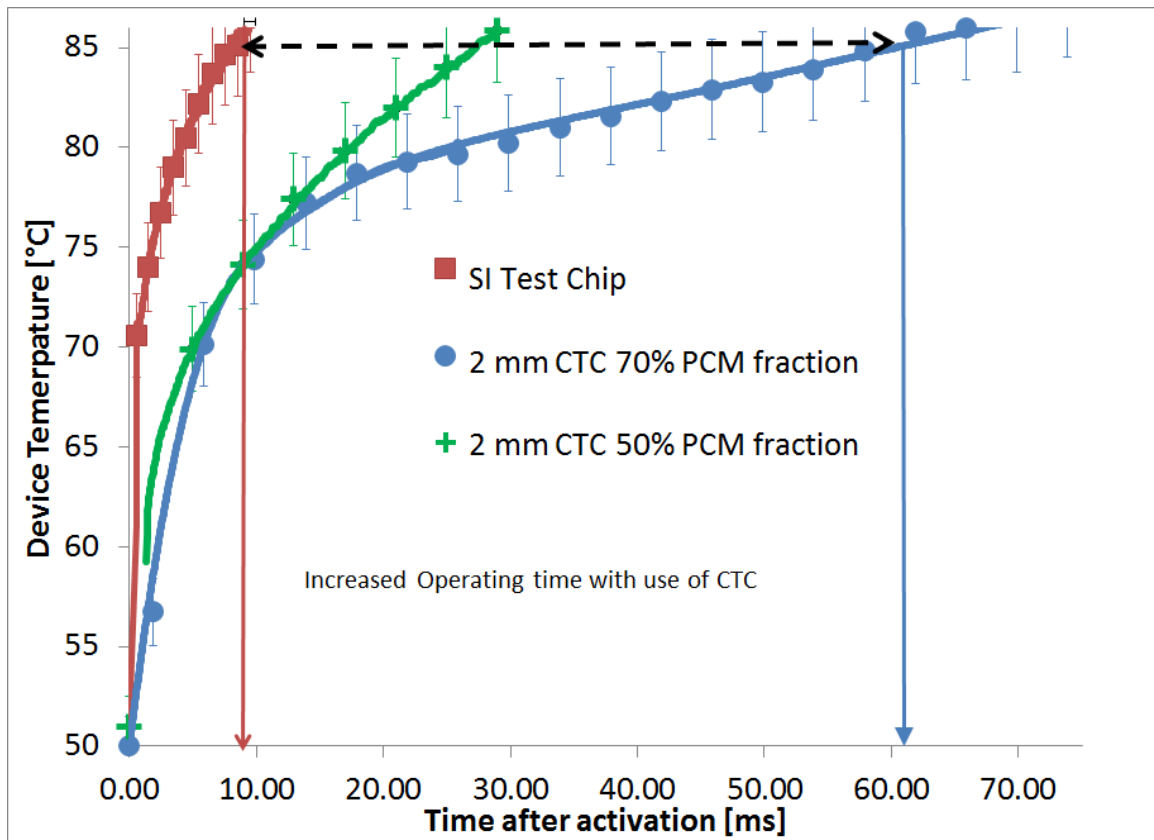


Figure 5.15 Temperature histories of $\sim 395\text{W}/\text{cm}^2$ heat flux hotspots with 2 mm diameter CTCs monolithically integrated as a part of the device under test, both compared to a Si baseline

Increasing the diameter of the CTC to 3 mm improved the performance at 50% PCM fraction (Figure 5.16). The allowable device operating time with a 50% CTC improves from 29 ms for the 2 mm device to 55 ms with the 3 mm device. This suggests that at 50% PCM fraction, the amount of PCM available for thermal storage may be less than ideal in the 2 mm CTC. Increasing the CTC diameter to 3 mm while keeping the PCM fraction constant provides more PCM for thermal storage.

However, the phenomena cannot be explained completely by the total PCM available for storage. The 2 mm 70% PCM fraction CTC outperforms both the 2 mm and 3 mm 50% CTC, even though the 3 mm CTC has more PCM available than the 70%/

2mm device (0.67 mm^3 vs 0.47 mm^3 , respectively). Another important input is the location of the PCM within the CTC. As seen in Figure 5.6, the 70% PCM fraction device has more PCM available near the center of the device, while the 50% PCM fraction devices have most of their PCM located near the outer edge of the PCM. If the melt front does not reach the edge of the CTC before reaching the throttle temperature, the device with more PCM in the region inside of the melt front (here the 70%-2mm CTC) will perform better.

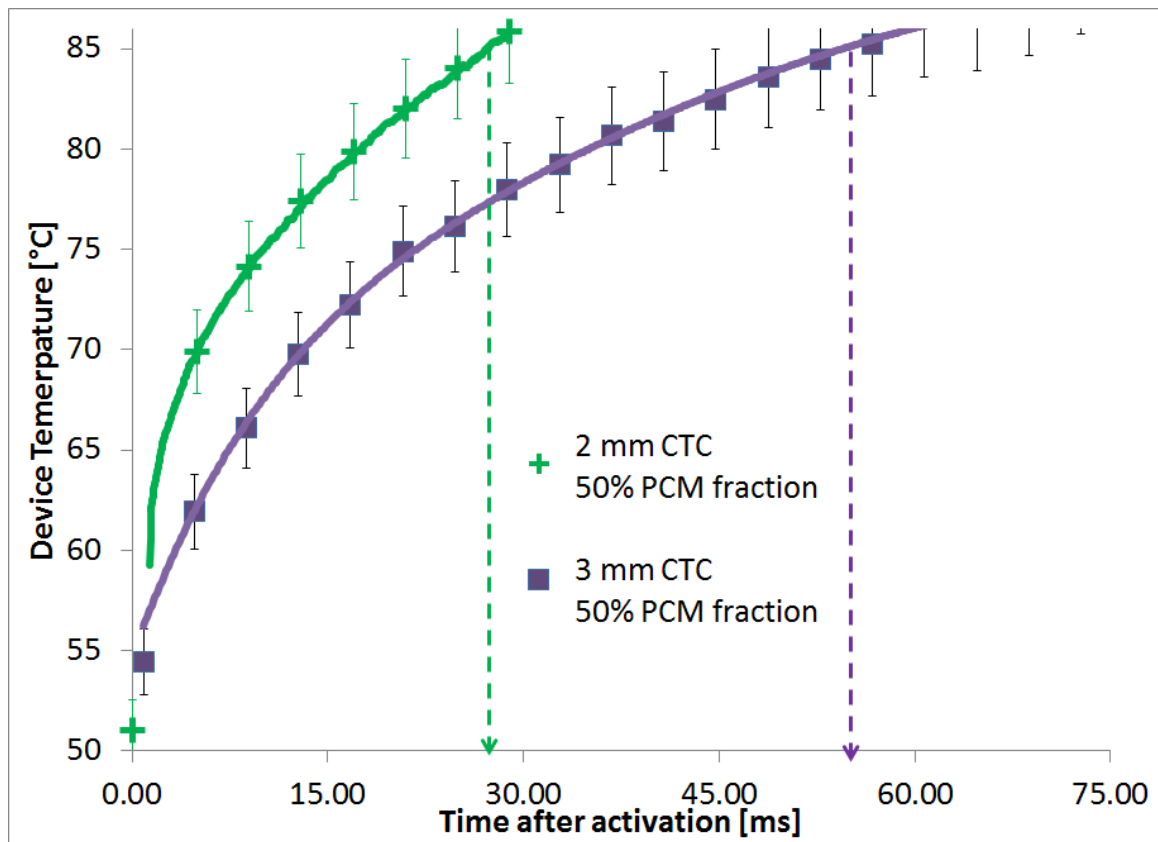


Figure 5.16 Comparison of temperature histories of $\sim 395 \text{ W/cm}^2$ heat flux hotspots with 2 mm and 3 mm 50% PCM fraction CTCs monolithically integrated as a part of the device under test at 395 W/cm^2

The notion that the PCM is not fully melted before reaching the throttle temperature is further supported by the device performance at 300 W/cm^2 , shown in Figure 5.17. At lower fluxes, the melt front will be able to travel further before the device reaches its throttle temperature, and yet the 3 mm-50% CTC loses its advantage over the 2 mm-50% CTC at this heat flux. This suggests that the available PCM in the 50%-2mm CTC was not exhausted prior to reaching its throttle temperature at 395 W/cm^2 . The melt front travels further, and yet the performance of the smaller diameter CTC gains ground on the larger diameter CTC. Instead, it appears that the smaller device is now able to access more of the PCM located at the periphery of the device prior to reaching its throttle temperature at this lower heat flux.

The 60%- 2 mm PCM fraction device provides the longest observed achievable operating time before reaching T_{throttle} at 300 W/cm^2 , at 101 ms. While this is significantly longer than the operating times observed at 395 W/cm^2 the relative improvement of the CTC over the Si baseline decreases to 4.2X at 300 W/cm^2 . This is because as the heat flux decreases, the average steepness of the temperature history for the Si baseline also decreases. As the throttle temperature gets closer to the steady state temperature of the device at lower heat fluxes, the rate of change of the temperature will also decrease in the latter portion of the transient. However, even at this lower heat flux, a 4.2X increase in allowable operating time is still a significant enhancement.




Figure 5.17 Temperature histories of $\sim 300\text{W}/\text{cm}^2$ heat flux hotspots with CTCs monolithically integrated as a part of the device under test, compared to a Si baseline

Even larger heat fluxes could potentially be addressed with the monolithic CTC than have been attempted at this time, depending on the desired operating times of the target high heat flux device. As discussed above, at higher heat fluxes the relative improvements realized from implementing the CTC concept (improvement over a Si baseline) should improve. Addressing larger heat fluxes was not attempted in this investigation because without direct indication of temperature during testing (voltage is measured directly and converted to temperature in post processing) it was difficult to ensure that the devices would not be burned or damaged during the testing process. Also, fabrication defects in the micro heaters that would result in localized increases in device

resistivity and temperature may not be evident in the measurement of total heater resistance until damage occurred.

5.4.3 Regenerative Cooling

Figure 5.18(a) shows the impact of regeneration, both from spreading into the substrate, as well as localized cooling from a TEC, on the overall system duty cycle. A significant reduction in the required cool down time of the test chip was achieved with the addition of TEC cooling (Figure 5.18(b)). Duty cycles of ~56% could be achieved with TEC cooling, compared to the ~43% achieved with regeneration through spreading into the substrate alone.

The bulk TEC used in this initial set of experiments was chosen because of its ease of implementation and availability, however the SSC's that would likely be used in an actual packaged CTC+SSC device would be thin film, microfabricated devices like those described in [8, 9, 25, 58]. Those thin film SSCs can dissipate heat fluxes 5-10 times larger than what the TEC used here delivered. Along with those higher performance SSCs would certainly come even better improvements in achievable device duty cycles.

That the TEC was able to deliver a noticeable improvement in device duty cycle is significant, because the 43% duty cycle achieved with the TEC off was due in large part to the liquid jet that was implemented to cool the TEC hot junction. Even with the TEC off and the additional 16 °C/W thermal resistance added to the thermal regeneration path due to the presence of the TEC [59], the regeneration is still more efficient by at

least an order of magnitude than the fan cooled approach, which will be discussed subsequently. In a packaged 3D device, it may be difficult to deliver this efficient liquid cooling to devices within the stack, however a SSC could be placed into proximity to the device while allowing its hot electrode to be located and cooled remotely [58, 67].




Figure 5.18 Duty cycle of test chips with CTC and attached TEC for regeneration (a) full cycle (b) close up of regeneration portion of cycle

As shown in Figure 5.19, activating the TEC just prior to throttling the device does not provide a significant reduction in the required regeneration time. The duty cycle of the device is nearly identical when the TEC is activated simultaneously with shutting the device down as it is when activating the TEC precedes shutting down the device by 50 ms. When the device is near its throttle temperature, it already has a significant temperature gradient between itself, the rest of the chip and the ambient to drive the heat out of the CTC. This is the reason for the rapid decrease in temperature just after the device is shut down. When the TEC is active during this part of the transient it provides little reduction in the required regeneration time. For regeneration, the TEC is most needed and should be operated during the portion of the transient when the temperature gradient between the device and its surroundings is lowest- during solidification and subcooling. This is not an obvious result, as thermoelectric devices have higher cooling powers at elevated temperatures. Taking advantage of the marginally improved cooling power at higher temperatures does not pay significant dividends in terms of regeneration rate however, because the device can already rapidly and efficiently reject heat at higher temperatures via spreading.

While operating the TEC continuously, or during some portion of the cycle prior to the throttle time has little impact on the required regeneration time, operating in this manner may still have potential for increasing duty cycles by extending the time the device can operate before requiring throttling. Increasing operating time while maintaining the same regeneration times would still have the effect of increasing the overall duty cycle. However, the effect of the regeneration method on the achievable device operating times was not studied as a part of this work.




Figure 5.19 Impact of TEC activation time on overall duty cycle

The large disparity in duty cycle between when the TEC is soldered to the CTC (Figure 5.8(b)) and when the TEC was soldered to the surface of the housing containing the water jet with no TIM between the TEC and the CTC (Figure 5.8(a)) is evident in Figure 5.20. The additional contact resistance between the cold junction and CTC in the latter configuration results in a degradation in the regeneration path that decreases the duty cycle from 43% to the 15% duty cycle observed in Figure 5.20. While this configuration (TEC with poor thermal coupling) may not be used in a production device, it gives a first indication of how a high thermal resistance regeneration approach negatively impacts the system duty cycles.




Figure 5.20 Temperature history of device with contact resistance at the TEC-CTC interface

When fan cooled regeneration was implemented similarly low duty cycles were observed, as shown in Figure 5.21. With nothing more than a fan blowing at the backside of the CTC, a 17% duty cycle was observed. While this is likely too low for implementation in a many core processor application, it is simple to implement, and would be more than sufficient for the low duty cycle applications such as high power RF devices discussed previously.




Figure 5.21 Temperature history of fan regenerated CTC

The liquid jet with no TEC configuration showed the most promise for high duty cycle application. With the TEC removed from the regeneration path, a duty cycle of 63% was achieved with jet impingement cooling (Figure 5.22). Also evident in Figure 5.22 is that the liquid cooled device is nearing its steady state temperature when it is throttled. The CTC with direct liquid cooled regeneration has the smallest effective RC time constant of the configuration considered. This is owed partially to the removal of the thermal mass of the TEC when compared to the TEC regenerated configuration. In addition, the removal of the TEC also removes ~ 16 °C/W of thermal resistance from the regeneration path, further reducing the effective RC time constant of the direct liquid jet cooled configuration.

Because the direct liquid jet cooled configuration is operating closer to its steady state temperature there is a significant decrease in the rate of increase in device temperature near the end of the pulse. If, for example the throttle temperature was slightly higher or if the cooling solution was slightly more efficient, a significant increase in device operating time could be had with little impact on the regeneration time. As a result duty cycles even higher than those observed here could be achieved.




Figure 5.22 Temperature history of liquid jet regenerated CTC

CHAPTER 6: COMPARISON OF MODEL PREDICTIONS TO EXPERIMENT

The numerical model formulation that has been implemented throughout this study can be compared the experimental data that has been collected. In order to test the model, the predictions of the 2D modeling assumption and numerical formulation described in chapter 3 will be compared to the device operating times observed through experiment.

6.1 Problem Formulation

In chapter 3 the potential for diamond shaped loops of high thermal conductivity materials of varying thicknesses filled with a hypothetical Bi-In type PCM was investigated in order to begin to understand the potential design space for CTC application. Based on the positive results of that investigation experimental prototypes described in chapter 5 were fabricated and tested with alloy-136 as the PCM and diamond shaped Si loops of 2 or 3 mm in diameter as the high thermal conductivity material. In order to quantify the predictive capacity of the numerical model, it has been applied to a computational domain designed to match the properties and dimensions of the experimental prototype more exactly.

The computational domain used for model validation takes advantage of the 45° symmetry of the problem just as was done in chapter 3. Two separate numerical domains were constructed. The first, shown schematically in Figure 6.1(a), mirrors the layout and dimensions of the 2 mm- 50% CTC prototype described in chapter 5.2. The second

computational domain, shown schematically in Figure 6.1(b), was constructed with dimensions matching the 2 mm -70% PCM CTC prototype described in chapter 5.2.

Unlike the models used in chapter 3, the mesh was extended beyond the outer diameter of the CTC. The 2D models used in chapter 3 had an adiabatic boundary condition applied at the outer edge of the 3 mm diameter CTC. In the experimental prototype however, there was only a single CTC and hotspot in a 2 cm square wafer. In order to more closely match the experimental conditions without excessive computational cost, the edges of the domain were extended to 6.7 mm from the center of the hotspot, where an adiabatic boundary condition was again applied.

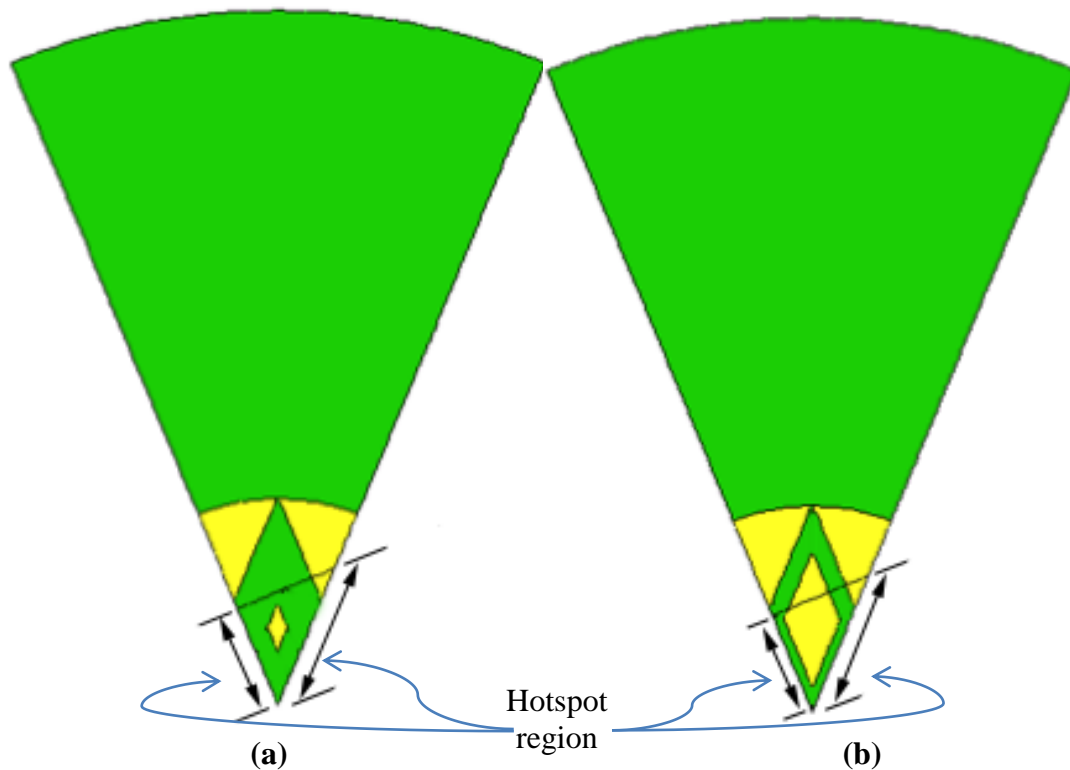


Figure 6.1 Schematic of computational domain for CTC model validation (a) 50% PCM fraction, (b) 70% PCM fraction

A $1.316 \times 10^{10} \text{ W/m}^3$ source term was applied in the hotspot region to simulate the heat generation of a 395 W/cm^2 hotspot spread over the $300 \text{ }\mu\text{m}$ thickness of the prototype wafer.

The numerical model only simulates the two dimensions that are in plane with the CTC, assuming no changes in the out of plane direction. However, the prototype device does have some variation in the normal direction due to the $110 \text{ }\mu\text{m}$ thick Si base that the $190 \text{ }\mu\text{m}$ thick CTC sits on. In order to account for the energy storage in the Si base as well as the CTC, effective values of specific heat, $c_{p,eff}$ and latent heat, $h_{sl,eff}$, of the PCM were defined using a rule of mixtures approach [54]. The effective properties were defined as:

$$c_{p,eff} = \frac{\rho_{PCM}c_{p,PCM}\left(\frac{z_{CTC}}{z_{Chip}}\right) + \rho_{Si}c_{p,Si}\left(\frac{z_{Si}}{z_{Chip}}\right)}{\rho_{PCM}} \quad (6.1)$$

$$h_{sl,eff} = h_{sl,PCM} \times \left(\frac{z_{CTC}}{z_{Chip}}\right) \quad (6.2)$$

Here, z_{CTC} , z_{Si} , and z_{Chip} represent the depth of the CTC, the thickness of the Si base that supports the CTC, and the total thickness of the test chip, respectively (Figure 6.2). The unmodified properties of alloy 136 used in the simulation were: $\rho = 8580 \text{ J/kg}$, $c_p = 133 \text{ J/kg} \cdot ^\circ\text{C}$, $\kappa = 10 \text{ W/m} \cdot ^\circ\text{C}$, and $h_{sl} = 18608 \text{ J/kg}$,




Figure 6.2 Schematic of a prototype CTC showing the relevant dimensions used in the numerical simulation

In addition to the energy stored in the Si, there will also be some enhancement in the thermal spreading of the hotspot due to the 100 μm thick Si substrate. In order to understand how this spreading enhancement due to the Si substrate affects the accuracy of the predictions of the 2D model, the problem was solved using κ_{PCM} values that both took into account and neglected the influence of the Si substrate. In one case, the value of κ_{PCM} used in the simulation was $10 \text{ W/m} \cdot ^\circ\text{C}$, which corresponds to the thermal conductivity of pure alloy-136. A second case was then considered where the impact of the Si substrate was considered by implementing an effective κ_{PCM} defined as:

$$\kappa_{PCM,eff} = \kappa_{PCM} \left(\frac{z_{CTC}}{z_{Chip}} \right) + \kappa_{Si} \left(\frac{z_{Si}}{z_{Chip}} \right) \quad (6.3)$$

6.2 Results

As seen in Figure 6.3, at 50% PCM fraction the 2D model prediction of when the average hotspot temperature reaches $85 \text{ }^\circ\text{C}$ agrees with the observed values from the

experiment. While the time to reach 85 °C is similar between model and experiment, there is some disagreement in the dynamics of the temperature rise early in the transient. The numerical model underpredicts the rate of hotspot temperature rise early in the transient because the 2D assumption is not fully valid during this time. When the hotspot is activated, the melt front must propagate vertically across the thickness of the CTC, a process that is not captured in a 2D model. Once the PCM directly above the hotspot is melted, the propagation of the melt front is primarily in the lateral (in-plane) direction and the predictions of the 2D model and the observations from the experiment begin to converge.




Figure 6.3 Comparison of model vs experiment for 50% PCM fraction CTC

At 70% PCM fraction, the 2D model under predicts the performance of the CTC when a κ_{PCM} value of $10 \text{ W/m} \cdot ^\circ\text{C}$ (the thermal conductivity of pure alloy-136) is used (Figure 6.4). Early in the transient, the model and experiment have similar trajectories, with the same temperature underprediction in the model due to the melt front propagation vertically across the CTC that was described for the 50% PCM fraction case. However as the melt front begins to propagate away from the hotspot after about 20 ms, the model and experiment begin to diverge.




Figure 6.4 Comparison of model vs experiment for 70% PCM fraction CTC

The 50% PCM fraction arrangement, unlike the 70% device, is already a reasonably efficient conductor/spreader due to the large amount of Si distributed in the

CTC matrix. As a result it is not strongly influenced by the 110 μm Si base that the CTC sits on. The Si in the 70% PCM fraction device, on the other hand, would not have been nearly as effective of a spreader without the Si base that the CTC pocket sits upon. If the 70% device were truly 2D, its performance would suffer, following the trajectory of the model prediction with $\kappa_{\text{PCM}} = 10 \text{ W/m} \cdot ^\circ\text{C}$.

Because of the presence of the Si base, the 70% PCM fraction device is able to take advantage of the large amount of PCM available for thermal storage in the CTC with an enhanced thermal conductivity that would otherwise require sacrifice of PCM fraction to achieve. When using the modified $\kappa_{\text{PCM},\text{eff}}$ (equation (6.3)) described above to capture the impact of the Si base on the melt front propagation, the model's prediction of the time to reach 85 $^\circ\text{C}$ matches better with the experiment. The evolution of the hotspot temperature during the transient is not captured by using the $\kappa_{\text{PCM},\text{eff}}$ approximation however. This is because thermal conductivity of the Si-PCM composite structure is anisotropic in nature and cannot be fully captured by a simple mixing rule. The ultimate enhancement in melt front propagation at the end of the transient due to the Si base is demonstrated by using the $\kappa_{\text{PCM},\text{eff}}$ approximation, however the approximation has little predictive ability to estimate the evolution of hotspot temperature earlier in the transient.

In general, the model is most accurate and effective when the 2D assumption holds. As a tool for exploring the potential CTC design space, the 2D numerical model maintains its usefulness. However, where highly accurate predictions of the chip temperature's temporal evolution are required, (for example when designing a control

scheme for the on-chip scheduler), a more detailed 3D or alternative model may be required.

CHAPTER 7: CONCLUSIONS AND FUTURE WORK

7.1 Summary

A new thermal management approach based on an imbedded PCM-high thermal conductivity material network (called a CTC) has been introduced that can be applied to dynamically operated microprocessors experiencing non-uniformities in their power profiles due to asymmetric architectures or integration of heterogeneous devices. The system has been analyzed from first principles using a reduced order model to identify the important parameters that drive achievable operating times and system duty cycles. A numerical model has been developed to study and optimize the arrangement of high thermal conductivity materials within a PCM melt to maximize the ability of a high heat flux device to store its energy efficiently in a CTC during its operating time slice. A 3D numerical model has been developed to study the performance of a CTC embedded and packaged within a 3D chip stack. The ability of solid state cooling (SSC) devices to rapidly regenerate the CTC has been investigated within the confines of the same 3D stack to highlight means of maximizing device duty cycles.

Experimental prototypes have been fabricated and tested to demonstrate the ability of a CTC to extend device operating times at heat fluxes up to 395 W/cm^2 . The results of the experimental investigation have been used to validate the numerical modeling used in the design of the system. Various means of regeneration of the CTC have been investigated experimentally, including air, liquid, and solid state cooling.

7.2 Conclusions and Key Findings

It has been shown that with careful design CTC integration can have a dramatic impact on core hopping and throttling frequencies, resulting in devices with greater spatial and temporal synchronization. One of the most important advantages of the approach is that it is locally passive, requiring no additional fluidic routing to realize its benefits.

Thermal conductivity, density, and latent heat of solid to liquid phase change of the CTC have potentially dramatic impact on the ability of a high heat flux device to store energy through lateral spreading, although diminishing returns are realized as the device performance approaches the constant temperature limit. Without proper design, PCMs inserted on the periphery of a high heat flux device can not only fail to enhance device operating times, but can also actually degrade system performance if the PCM behaves like an insulator. One of the more important parameters in deciding on the layout of a PCM matrix is the allowable temperature budget. For small temperature budgets it may be difficult to drive the melt front any significant distance through spreading, while at larger temperature budgets the divergence in thermal capacity enhancement between poorly and well-designed PCMs becomes large.

Using diamond shaped loops as the basis for the high thermal conductivity spreader matrix can result in an up to 20X increase in the achievable device operating times when diamond is chosen as the spreader material. More moderate thermal conductivity materials such as copper result in an operating time enhancement of up to 10X.

If implementation of a separately attached CTC is desired instead of a monolithic integration, the TIM becomes an integral part of the design process and must be absolutely minimized in order to extract benefits from the device. For traditional spreadable interface materials, the TIM and CTC should be kept under pressure to improve the ability of the electrical device to rapidly deposit energy into the CTC.

Devices that use Si as the spreader material in a diamond shaped loop configuration were fabricated and tested, demonstrating operating time enhancements of up to 7X at 395 W/cm^2 . The potential operating time enhancements achievable with CTC increase with increasing heat flux dissipated.

For CTC arrangements where the PCM fraction is not radially uniform, the appropriate metric for the CTC design should be based on the PCM fraction only in the region where the melt front reaches before the throttle temperature is reached.

Increasing the throttle temperatures of the devices to levels closer to their steady state temperatures can significantly increase device duty cycles by taking advantage of the change in the rate of increase in temperature at higher temperatures.

The addition of active regeneration dramatically increases the system availability over what can be achieved with more passive regeneration approaches, although still falling short of the availability that would be seen with the unaltered silicon substrate. However, devices with CTC integration can take advantage of improvements in duty cycle through device redundancy that the unaltered Si devices may not be able to utilize efficiently.

SSC regeneration can offer a 2X-3X reduction in required regeneration times compared with spreading into the substrate alone. Liquid cooling is the most efficient and

effective means of regenerating the CTC and should be implemented where available and practical.

When SSC regeneration is used, the cooler should be operated simultaneous to or just after the throttle event, as the devices are most helpful when then CTC has cooled somewhat and have little temperature gradient left to drive heat out of the device.

Duty cycles of 15%-20% can be achieved with no additional active regeneration, while achieving duty cycles in excess of 60% may require a highly efficient CTC regeneration approach.

7.3 Original Contributions

1. Conceptual development of a hybrid (heat spreading and solid-liquid phase change) thermal management strategy enabling composite thermal capacitors.
2. Critical comparison of a set of promising CTC topologies (interconnected islands and diamond shaped loops) and quantification of key design trade-offs- PCM loading, effective properties, temperature budget, heat flux- using reduced order models.
3. Explored the use of asymmetric cooling strategies for the CTC using thermoelectric cooling, and liquid and air convective cooling. Developed recommendations for optimal methods of regeneration for different potential applications.
4. Developed a CMOS- compatible fabrication process that allowed implementation of the CTC family of devices with integrated thermometry capabilities to enable a comprehensive performance characterization of the CTC in an environment relevant to practical application.

5. A new experimental capability for measuring rapid thermal transients with [ms] time scales using resistive thermometry has been developed which has broad applicability for thermal characterization

Suggested Future Work

1. An integral piece of getting the CTC concept implemented in a production microprocessor environment is a more rigorous quantification of the computational impact of the thermal approach. As a result, a crucial next step is a combined electrical- thermal analysis that uses the thermal performance benefits of the CTC as an input to determine what additional computational benefits can be extracted from a system that has implanted the CTC.
2. A numerical model has been built and validated to allow the investigation and analysis of a number of important parameters that may impact the performance of a CTC that has been integrated into a 3D stack. As the concept moves closer to being implemented, the model should be exercised to allow the designer to tune the parameters of the system to most efficiently use the state of the art materials at the time (e.g. inter-layer bonding method improvement, increases in SSC ZT, etc.)
3. Enhancements to lateral spreading have been studied extensively as a part of this thesis; however enhancements in vertical (across the device) thermal conductivity may also play an important role in maximizing system availability by improving (speeding up) CTC regeneration. An analysis should be conducted to determine the important tradeoffs between in and out of plane thermal conductivity enhancement in order to better optimize the design of the overall system.

4. The experimental prototypes that were tested as a part of this investigation all used Si as the spreader material; however the models predict that significant enhancements above what have been demonstrated here can be achieved with the use of higher thermal conductivity spreader materials. Development of a process that allows easy integration of some of these higher κ materials into a production CTC would be a worthwhile endeavor.
5. The class of PCM that was focused on in this investigation is metallic and thus electrically conductive. An investigation should be conducted to quantify what, if any impacts embedding electrically conductive materials into the substrate may have on the device performance (e.g. parasitic capacitive coupling or ground plane interference)

APPENDIX A: EXPERIMENTAL DEVICE FABRICATION

A.1 Pyrex Test Chip

The Pyrex test chip described in section 5.1 was fabricated on a 500 μm thick polished Pyrex wafer. The process flow for the Pyrex test chip fabrication is shown in Figure A.1. Prior to performing any microfabrication steps, the wafer was scrubbed using an acetone saturated swab, followed by a 20 minute 3:1 H_2SO_4 : H_2O_2 piranha clean at 120 $^\circ\text{C}$. After cleaning the wafer was dried and dehydrated in a 150 $^\circ\text{C}$ oven for 2.5 hours.

After these wafer preparation steps were completed, MicroPrime MP-P20 hexamethyldisilazane (HMDS) was applied to the wafer using a spin coater at 3000 RPM for 30 s. The HMDS improves adhesion of the photoresist which will be applied in the next process step by removing any adsorbed water and lowering the surface energy [68]. After application, the HMDS was baked for 60s on a 100 $^\circ\text{C}$ hotplate for 60s.

NR9-1500PY negative photoresist is then spun onto the wafer for 40s at 3000 RPM. This should result in a 1425-1575 nm thick coating of the photoresist. After application the soft bake of the photoresist is performed on a 150 $^\circ\text{C}$ hotplate for 3.5 minutes. The soft bake serves to drive off some of the solvents in the photoresist, preventing the wafer from sticking to the mask in subsequent steps. The soft bake also improves adhesion of the resist to the wafer and the stability of the resist during further processing steps.

The wafer is then aligned with the mask and exposed to a 365 nm light source for a total target exposure of 285 mJ/cm^2 . The photo exposure is performed using a mercury

lamp that exposes across the 230 nm to 405 nm spectrum. The alignment and exposure tool (a Karl Suss MA-6 mask aligner) controls the intensity of the lamp such that its output at the 365 nm wavelength is constant throughout the exposure. Thus the 285 mJ/cm² exposure is computed only based on the irradiation at 365 nm, even though exposure at other wavelengths is occurring simultaneously. Just as is done here, photo exposure doses stated throughout this dissertation consider only the dose at the target wavelength, and neglect radiation at all other wavelengths.

After exposure, a post exposure bake is performed on a 100°C hotplate for 3.5 minutes. The resist is then developed in RD6 developer for 12 s, using gentle agitation of the developer bath. Small amounts of resist may still be present on the surface of the wafer, even after developing, so a short 15s descum process is performed prior to perming any subsequent steps to prevent defects from occurring in any deposited films. The descum is performed in a Plasma-Therm (PT) Reactive Ion Etcher (RIE), which uses a 13.56 MHz plasma to etch the photoresist isotropically. The descum process uses a 100 mT O₂ plasma at 90 sccm. 300 W of RF power is applied during the process.

After the resist patterning process is complete the Pt film that will serve as the heaters and RTD is deposited with Ti adhesion layer using Electron beam (E-beam) evaporation. A 250 Å Ti layer is first deposited at a rate of 1 Å/s to allow the Pt film to adhere to the Pyrex wafer. Next, a 2500 Å thick Pt film is deposited on top of the Ti layer at a rate of 3 Å/s. The high temperature of the Pt during evaporation tends to crack or warp the photoresist during long deposition periods so the full 2500 Å is not deposited in one step. Instead, after each 1000 Å of deposition, the electron gun is turned off and the

wafer is allowed to cool for 5 minutes prior to resuming deposition. The E-beam chamber vacuum is maintained throughout this entire process.

After the metal deposition, lift off is performed using Microposit Remover 1165a resist stripper. To enhance the resist stripper's ability to completely remove the photoresist from the wafer, ultrasonic agitation and 50 °C heating was applied to the bath.

In the next step a new layer of negative photoresist will be spun onto the wafer to allow for the patterning of the heater and RTD leads. Just as before HMDS is applied first to promote adhesion. The thicker, NR71-3000P is chosen as the photoresist for this step to ensure that the sidewalls are tall enough after patterning and subsequent metallization to allow for a successful lift-off process. The NR71-3000P is applied at 3000 RPM for 40 s, which should result in a ~2850-3150 nm thick resist layer.

A 4 minute, 165 °C hotplate soft bake is next performed to prepare the wafer and resist for photolithography. The wafer is aligned with the mask and exposed to 365 nm light for a total target dose of 60.9 mJ/cm². After exposure, a post exposure bake is performed at 110°C in an oven. A hotplate is not used for the post exposure bake, because the low κ Pyrex substrate makes it somewhat difficult to get an satisfactory curing of the entire photoresist film when heating is only supplied on the one side of the wafer. When baking, the wafer should be supported above the oven shelf using supports that only touch the wafer's edges. Using backside supports that touch the wafer in regions that correspond to the patterns in the photoresist can result in nonuniform heating in these crucial regions.

After post exposure bake, the resist is then developed in RD6 developer for 60s, using gentle agitation of the developer bath. After developing, a 30 s descum is performed to ensure that the substrate is fully clean prior to metallization.

Next, the metallization for the device leads is performed using E-beam evaporation. A 500 Å Ti layer is deposited, immediately followed by a 4850 Å layer of Au. Lift-off is then performed in 1165a resist stripper at 60 °C.

In order to passivate the metal elements that have been deposited on the substrate (the heaters and RTDs) and protect the devices from exposure to the PCMs, a 2 μm thick conformal layer of Si₃N₄ is deposited on top of the devices using a Uniaxis Plasma Enhanced Chemical Vapor Deposition (PECVD) system. The deposition process uses SiH₄, NH₃, and N₂ as precursor gases that form the Si₃N₄ in a 13.56 Mhz plasma, which increases the deposited film's ability to conform to the topology of the features on the surface of the wafer. He dilution is added to the gas mixture to alleviate stress in the resulting film [69]. The parameters implemented in the Si₃N₄ deposition are detailed in Table A.1.

After the Si₃N₄ deposition, a final photolithography step is performed to open small windows in the film to allow for wire bonding. Shipley 1827 positive photoresist is spun onto the wafer at 3000 RPM for 60s. A 90s soft bake is performed in a 115 °C oven to prepare the resist for photo processing. The wafer is then aligned with the mask and exposed to 405nm light for a total dose of 300 mJ/cm².

The 1827 photoresist is then developed in Microposit MF-315 developer for 70s followed by a 20s descum process. The patterned resist is then baked for 30 minutes at

110 °C to harden it and improve its resistance to the Si₃N₄ plasma etching that will follow.

Table A.1 Si₃N₄ deposition process parameters

| | |
|--------------------------------------|----------|
| Temperature | 250 °C |
| Pressure | 1100 mT |
| SiH ₄ (5% in He) Gas Flow | 200 sccm |
| NH ₃ Gas Flow | 8 sccm |
| He Gas Flow | 560 sccm |
| N ₂ Gas Flow | 150 sccm |
| Power | 50 W |

The final step in the Pyrex test chip fabrication process is the opening of the bond pad windows in the Si₃N₄ film using a RIE process. The film is bombarded with a reactive RF plasma, which selectively etches the Si₃N₄ while leaving the photoresist intact. As a result, square windows are opened at the tips of the heater and RTD leads that allow the devices to be wire bonded to an external circuit board. The process recipe for the Si₃N₄ etch is listed in Table A.2.

Table A.2 Si₃N₄ RIE etch process parameters

| | |
|------------------|---------|
| Temperature | 34 °C |
| Pressure | 45 mT |
| CHF ₃ | 45 sccm |
| O ₂ | 5 sccm |
| Power | 350 W |




Figure A.1 Process flow diagram for Pyrex test chip fabrication

A.2 Standalone CTC fabrication

The CTC that was used in the separately attached CTC experiments described in section 5.1 was fabricated by etching the patterned cavities into a 190 μm thick Si wafer using an Inductively Coupled Plasma (ICP) RIE process. The major CTC fabrication steps are shown schematically in Figure A.2.

The Si wafer was prepared using the same acetone scrub and piranha clean process that was described in section A.1. Next, a 2.75 μm thick film of SiO_2 was deposited on the wafer using the Uniaxial PECVD. The SiO_2 will serve as a mask during the Si etching process that will follow. SiO_2 is chosen because it has a superior selectivity to the Si ICP etch when compared to photoresist. The SiO_2 mask will be patterned using a photolithographic process.

Because deep trenches will eventually be etched in the CTC wafer a second 500 μm thick carrier wafer is mounted to the back side of the CTC wafer using Cool Grease 7016 vacuum grease. This vacuum grease is chosen because its relatively high thermal conductivity (3 $\text{W}/\text{m}\cdot^\circ\text{C}$) will allow better temperature control of the CTC wafer during the ICP etch process [70]. Wafer cooling is primarily provided via He gas blown on the backside of the wafer during the etch, so minimizing the thermal resistance between the carrier wafer and the CTC wafer helps to maintain tighter control over the process conditions.

After mounting, Shipley 1827 positive photoresist is spun onto the wafer at 3000 RPM for 60s. A 60s hotplate soft bake is then performed to prepare the wafer for

photolithography. The wafer is then aligned with the mask and the 1827 photoresist is patterned and baked using the same process detailed in section A.1.

After the photoresist is patterned the wafer is placed in PT-ICP tool where the photoresist will be used as a mask to etch the SiO₂ layer. The PT-ICP using a ICP-RIE process, which differs from standard RIE processes such as that described in section A.1 in that the PT-ICP uses a second RF power supply that allows a significantly more dense plasma to be generated when compared to the PT-RIE tool. This dense plasma enables anisotropic etching in the PT-ICP, unlike the PT-RIE which etches nearly isotropically. The recipe used to etch the SiO₂ mask is shown in Table A.3

Table A.3 PT-ICP oxide mask etch parameters

| | |
|---------------------------------------|---------|
| Ar | 5 sccm |
| C ₄ F ₈ | 15 sccm |
| CO ₂ | 28 sccm |
| RF power 1 (plasma acceleration) | 40 W |
| RF power supply 2 (plasma generation) | 800 W |

After the SiO₂ mask is etched, the CTC cavities that will later be filled with PCM are etched into the wafer using a Bosch process in the PT-ICP. A Bosch process alternates between deposition of a passivation film and etching steps in order to enhance the achievable aspect ratios of the Si etch. The passivation film is a fluorinated polymer which is deposited on the surface using C₄F₈. Next a two part etching step is performed using SF₆. In the first part of the etch, the polymer on the bottom surface of the trench is

etched away at a faster rate than the polymer on the sidewalls. In the second part of the etch, the now exposed Si in the bottom of the trench is etched away, again using SF₆. This entire process takes about 10s and the cycle repeated continuously until a trench depth of 130 μm is achieved. The recipe for the Bosch process used to etch the CTC trenches is shown in Table A.4.

Table A.4 Bosch process parameters

| | Deposition | Etch A | Etch B |
|--|------------|--------|--------|
| Pressure [mT] | 15 | 15 | 15 |
| C ₄ F ₈ [sccm] | 100 | 0 | 0 |
| SF ₆ [sccm] | 0 | 100 | 150 |
| Ar [sccm] | 40 | 40 | 40 |
| RF power 1 (plasma acceleration) [W] | 0 | 12 | 10 |
| RF power supply 2 (plasma generation) [W] | 825 | 825 | 825 |
| Time [s] | 3 | 4 | 6 |

After the trenches are etched, any remaining SiO₂ is stripped using 6:1 buffered oxide etchant (BOE). Finally, a 1 μm thick Pt layer with 250 Å Ti adhesion layer is sputtered conformally onto the surface of the CTC cavities and sidewalls using the Unifilm sputterer.




Figure A.2 Process flow of the major stand-alone CTC fabrication steps

A.3 PCM selection and loading process

After the CTC Si spreader matrix has been fabricated it must be filled with the selected PCM. This process is followed for the CTCs described in sections 5.1 and 5.2.

Alloy 136 was chosen as the PCM to be used for experimental testing, due to its appropriate melting temperature (58 °C), as well as its relatively small change in volume upon phase change [71]. The Bismuth in alloy 136 causes it to expand upon solidification, unlike most materials, which contract. Compared to some other commercially available materials in its melting temperature range, alloy 136's expansion (or contraction) is small. Furthermore, alloy 136's coefficient of thermal expansion for single phase heating of $23 \times 10^{-6}/^{\circ}\text{C}$ [46] is closer to that of Si than some of the polymeric materials often used for wafer bonding and integrating microstructures into 3D stacks such as benzocyclobutene ($\sim 61 \times 10^{-6}/^{\circ}\text{C}$) [72] or SU-8 ($\sim 52 \times 10^{-6}/^{\circ}\text{C}$) [73]. This suggested that integrating the alloy-136 into a 3D stack may be feasible from a thermal expansion perspective.

While the thermal expansion characteristics of alloy 136 may be favorable, a potential downside of the alloy is that it is not particularly wetting [46]. The non-wetting behavior of the alloy can make manually filling the CTC cavities particularly difficult. In order to improve the wetting behavior of the PCM, the surface of the Si cavities was coated with Pt. Pt was chosen because as a noble metal under normal conditions it does not form appreciable surface oxides which lower the surface energy and reduce wetting. Unlike other noble metals such as Au or Ag, Pt also has an extremely low dissolution rate in most solders which is important to the long term stability of the PCM [55, 56].

The Pt coating does improve the PCM's affinity to wet the surface of the CTC, but some additional assistance is needed in order to ensure that the cavities have filled completely. For each device described in section 5.2, except the 2 mm 60% PCM fraction

CTC, the cavities were coated with flux prior to application of the PCM. Even though the surface is coated with Pt, surface oxides still exist on the surface of the PCM that the flux can remove, thus improving the surface wetting and cavity filling.

After Pt coating and optional flux application, the CTC was heated to slightly above the melting temperature of the PCM and a small amount of PCM was then placed on top of the CTC. The Pt coating allows the PCM to adhere loosely to the top surface of the CTC, but the PCM still will not fill the cavity below without some additional force.

To provide this additional forcing, the CTC is first removed from the heat so that the PCM on the surface can re-solidify. After the PCM has solidified, it is covered with polytetrafluoroethylene (PTFE) tape. The PTFE tape has nonstick properties that keep the PCM from sticking to any other materials used during the CTC filling process. The PTFE tape also creates an effective seal around the edges of the CTC pocket so that when force is applied to the PCM, it will flow into the CTC cavities instead of sliding on the top surface of the wafer.

The PCM is then reheated to just above its melting temperature, and a flat object such as a glass slide is placed on top of the molten PCM. Additional force is then applied to the surface of the glass slide, either manually or by placing a heavy object on top of the slide. The weight of the slide forces the molten PCM into the CTC cavities. After the weight is applied the entire structure is left heated for approximately five minutes.

Next, the heat is again removed from the CTC so that the PCM can re-solidify. Once the structure is cooled, the weight, glass slide and PTFE tape are removed, and the CTC is inspected to ensure that no visible voids or air cavities are present.

A.4 Monolithically integrated CTC and heater/RTD test chip

The monolithic test chip described in section 5.2 was fabricated using processes similar to those described in section A.1 for making the heaters and RTDs and section A.2 for making the CTC. The areas where the fabrication process differs from that already described will be detailed in this section. A process flow of the major fabrication steps for the monolithically integrated CTC and heater/RTD test chip is shown in Figure A.3.

A 500 μm thick Si wafer used as the substrate that the heaters, RTDs and CTC were built into. A full RCA clean was performed to prepare the wafer as opposed to the simple piranha clean that was described in the previous sections. This more exhaustive wafer cleaning process was performed because a gate oxide will eventually need to be grown on the wafer to protect the heaters and RTDs from being short circuited by the electrically conductive Si substrate. The RCA clean consists of three steps that combine to ensure that a quality, low defect gate oxide can be grown on the substrate. The first step removes any organic contaminants on the wafer. Next the native SiO_2 layer is stripped from the wafer. Finally any metallic contaminants remaining on the wafer are removed from the wafer in an ionic cleaning step. The steps undertaken in the RCA clean are detailed in Table A.5.

Table A.5 RCA clean parameters

| Clean Step | Mixture | Time [min] | Temperature [°C] |
|-------------|--|------------|------------------|
| Organic | 5:1:1 H ₂ O:H ₂ O ₂ :NH ₄ OH | 10 | 75 |
| Oxide Strip | 6:1 BOE | 0.25 | 25 |
| Ionic | 6:1:1 H ₂ O:H ₂ O ₂ :HCl | 10 | 75 |

After preparation, 500 nm thick SiO₂ film was grown on the Si wafer through Low pressure chemical vapor deposition (LPCVD) in a wet atmosphere in a 1000 °C, ambient pressure furnace. Next, the steps described above in section A.1 for fabricating the heaters and RTDs were followed to deposit those devices on top of the gate oxide. Because the Si wafer has a high thermal conductivity, the following soft bake and post-exposure bake procedure was followed (all steps using hotplate bakes):

Table A.6 Photoresist bake times for Si substrate

| | Soft bake temperature [°C] | Soft bake time [s] | Post bake temperature [°C] | Post bake time [s] |
|------------|----------------------------------|-----------------------|----------------------------------|-----------------------|
| NR9-1500PY | 150 | 60 | 100 | 60 |
| NR71-3000P | 150 | 60 | 110 | 60 |

The device leads, which were made of 4850 Å of Au on the Pyrex test chip, were replaced with a composite film of 2000 Å Au on top of 3000 Å Cu. The change was not

due to a performance deficiency in the completely Au leads. Instead, using the Cu/Au composite film is simply less expensive and time intensive to make.

In order to limit the number of different kinds of film deposited on the Si substrate, a 1 μm SiO_2 film was used as a passivation layer for the heaters and RTDs, as opposed to the Si_3N_4 film that was used on the Pyrex test chip. The SiO_2 film was deposited using the Uniaxial PECVD using the parameters in Table A.7.

Table A.7 SiO_2 PECVD deposition parameters

| | |
|-----------------------------|-----|
| Temperature [°C] | 250 |
| Pressure [mT] | 900 |
| SiH_4 [sccm] | 400 |
| N_2O [sccm] | 900 |
| Power [W] | 25 |

After the heaters and RTDs were fabricated and passivated a 1.5 μm film of PECVD SiO_2 was grown on the reverse side of the wafer using the parameters in Table A.7. Because 1 μm of LPCVD had already been deposited during the gate oxide growth, a total of 2.5 μm of SiO_2 is now present on the backside of the wafer. This oxide layer will serve as the mask for the CTC fabrication that will follow.

The photoresist mask for the etching the CTC cavities was processed and patterned using the same parameters that were described in section A.2. Backside alignment cameras were used to ensure that the CTC cavities were aligned with the heaters and RTDs on the opposite side. The mask was used to pattern the photoresist.

Next the wafer was mounted on a carrier wafer using Megaposit SPR 220 7.0. This thick film photoresist was used for mounting instead of the vacuum grease that was used in section A.2 because the processing temperatures as well as the high energy plasmas that the wafer are exposed to can cause charring in the thermal grease. This can make detachment of the carrier wafer as well as complete removal of any residue from the thermal grease difficult after processing. As a result, photoresist was used for mounting the carrier wafer to avoid damaging the heaters and RTDs that had already been fabricated on the reverse side of the wafer. Using photoresist does make temperature control during ICP etching more of a concern; however this was considered an acceptable trade off.

The SPR 220 is spun on the carrier wafer at 3000 RPM for 50s. The carrier wafer is then attached to the sample wafer and the SPR 220 is cured during the same 30 min, 110 °C bake used to cure the 1827 photoresist mask.

After mounting, the oxide mask was etched in the ICP using the parameters in Table A.3. Next, 190 μm deep CTC cavities were etched using the Bosch process described in Table A.4. After etching the cavities, the oxide mask is stripped using 6:1 BOE.

At this point there are 190 μm deep CTC cavities etched into a 500 μm thick wafer. Ideally the bottoms of the cavities should be as close as possible to the heaters on the reverse side of the wafer. To achieve this result, the entire wafer is now etched using as Bosch process with no oxide mask in place. The process parameters described in Table A.4 will etch the bottoms of the trenches faster than the top surface of the wafer. Because of this fact, the process pressure was reduced to 12 mT from the 15 mT described in

Table A.4 to ensure that the bottoms and tops of the trenches etched at the same rate. Using this process the wafer was thinned such that the final sample had 190 μm deep trenches in a 300 μm thick wafer.

After the CTC cavities have been etched, the final steps are: Pt coating of the cavities as described in section A.2, followed by opening of the wire bonding windows as described in section A.1. The RIE etch used to open the windows in the SiO_2 passivation layer used the parameters in Table A.8.

Table A.8 Si_3N_4 RIE Etch Process Parameters

| | |
|----------------|---------|
| Temperature | 34 °C |
| Pressure | 45 mT |
| CHF_3 | 45 sccm |
| O_2 | 5 sccm |
| Power | 350 W |

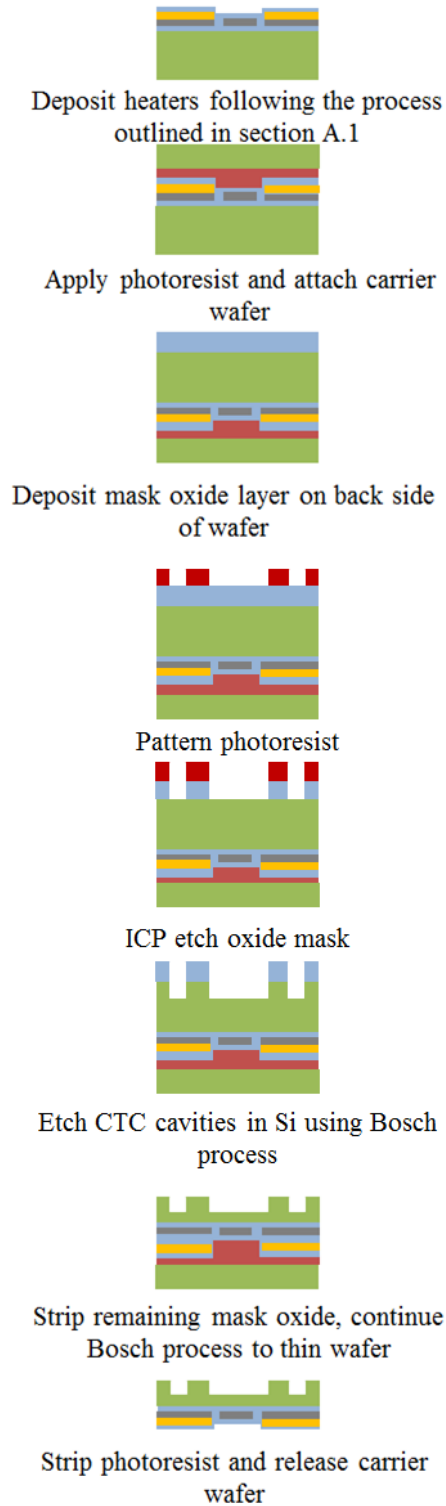


Figure A.3 Process flow of major fabrication steps for monolithically integrated CTC and heater/RTD test chip

REFERENCES

- [1] L. B. Kish, "End of Moore's law: thermal (noise) death of integration in micro and nano electronics," *Physics Letters A*, vol. 305, pp. 144-149, 2002.
- [2] C. Green, A. G. Fedorov, and Y. K. Joshi, "Fluid-to-Fluid Spot-to-Spreader (F2/S2) Hybrid Heat Sink for Integrated Chip-Level and Hot Spot-Level Thermal Management," *Journal of Electronic Packaging*, vol. 131, p. 025002, 2009.
- [3] R. Viswanath, W. Vijay, A. Watwe, and V. Lebonheur, "Thermal performance challenges from silicon to systems," *Intel Technology Journal*, vol. 3, pp. 1-16, 2000.
- [4] D. Koester, R. Venkatasubramanian, B. Conner, and G. J. A. S. G. J. Snyder, "Embedded thermoelectric coolers for semiconductor hot spot cooling," in *The Tenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems, IITHERM '06.* , 2006, pp. 491-496.
- [5] W. Peng and B.-C. Avram, "On-chip hot spot cooling using silicon thermoelectric microcoolers," *Journal of Applied Physics*, vol. 102, p. 034503, 2007.
- [6] N. Sridhar, K. Lofgreen, D. Chau, and G. Chrysler, "Thin film thermoelectric cooler thermal validation and product thermal performance estimation," in *The Tenth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronics Systems, IITHERM '06.* , 2006, pp. 470-475.
- [7] V. Sahu, Y. K. Joshi, and A. G. Fedorov, "Hybrid Solid State/Fluidic Cooling for Hot Spot Removal," *Nanoscale and Microscale Thermophysical Engineering*, vol. 13, pp. 135-150, 2009.

- [8] V. Sahu, Y. K. Joshi, and A. G. Fedorov, "Experimental investigation of hotspot removal using superlattice cooler," in *12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, ITherm 2010* 2010, pp. 1-8.
- [9] A. Shakouri, "Nanoscale thermal transport and microrefrigerators on a chip," *Proceedings of the IEEE*, vol. 94, pp. 1613-1638, 2006.
- [10] C. E. Green, A. G. Fedorov, and Y. K. Joshi, "Scaling Analysis of Performance Tradeoffs in Electronics Cooling," *Components and Packaging Technologies, IEEE Transactions on*, vol. 32, pp. 868-875, 2009.
- [11] S. Narayanan, A. G. Fedorov, and Y. K. Joshi, "On-chip thermal management of hotspots using a perspiration nanopatch," *Journal of Micromechanics and Microengineering*, vol. 20, p. 075010, 2010.
- [12] A. Bar-Cohen, M. Arik, and M. Ohadi, "Direct liquid cooling of high flux micro and nano electronic components," *Proceedings of the IEEE*, vol. 94, pp. 1549-1570, 2006.
- [13] A. W. Topol, D. C. La Tulipe Jr, L. Shi, D. J. Frank, K. Bernstein, S. E. Steen, A. Kumar, G. U. Singco, A. M. Young, and K. W. Guarini, "Three-dimensional integrated circuits," *IBM Journal of Research and Development*, vol. 50, pp. 491-506, 2006.
- [14] S. Narayanan, A. G. Fedorov, and Y. K. Joshi, "Gas-Assisted Thin-Film Evaporation from Confined Spaces for Dissipation of High Heat Fluxes," *Nanoscale and Microscale Thermophysical Engineering*, vol. 13, pp. 30-53, 2009.

- [15] R. Yavatkar and M. Tirumala, "Platform wide innovations to overcome thermal challenges," *Microelectronics Journal*, vol. 39, pp. 930-941, 2008.
- [16] S. Borkar, "Thousand core chips: a technology perspective," in *44th Annual ACM IEEE Design Automation Conference*, 2007, pp. 746-749.
- [17] M. D. Hill and M. R. Marty, "Amdahl's law in the multicore era," *Computer*, vol. 41, pp. 33-38, 2008.
- [18] W. Huang, K. Skadron, S. Gurumurthi, R. J. Ribando, and M. R. Stan, "Exploring the Thermal Impact on Manycore Processor Performance," in *26th IEEE SEMI-THERM Symposium*, 2010.
- [19] Z. Xiuyi, Y. Jun, X. Yi, Z. Youtao, and Z. Jianhua, "Thermal-Aware Task Scheduling for 3D Multicore Processors," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 21, pp. 60-71, 2009.
- [20] Y. J. Kim, Y. K. Joshi, A. G. Fedorov, Y.-J. Lee, and S.-K. Lim, "Thermal Characterization of Interlayer Microfluidic Cooling of Three-Dimensional Integrated Circuits With Nonuniform Heat Flux," *Journal of Heat Transfer*, vol. 132, pp. 0410091-0410099, 2010.
- [21] M. P. Gupta, M. Cho, S. Mukhopadhyay, and S. Kumar, "Thermal mangament of multicore processors using power multiplexing," in *12th IEEE Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems, IITHERM 2010* 2010, pp. 1-7.
- [22] P. Chaparro, J. González, G. Magklis, Q. Cai, and A. González, "Understanding the Thermal Implications of Multicore Architectures," *IEEE Transactions on Parallel and Distributed Systems*, vol. 18, pp. 1055-1065, 2007.

- [23] M. Cho, N. Sathe, M. Gupta, S. Kumar, S. Yalamanchilli, and S. Mukhopadhyay, "Proactive power migration to reduce maximum value and spatiotemporal non-uniformity of on-chip temperature distribution in homogeneous many-core processors," in *26th Annual IEEE SEMI-THERM Symposium*, 2010, pp. 180-186.
- [24] "Failure Mechanisms and Models for Semiconductor Devices," ed: JEDEC Solid State Technology Association, 2009.
- [25] A. Bar-Cohen and P. Wang, "On-chip Hot Spot Remediation with Miniaturized Thermoelectric Coolers," *Microgravity Science and Technology*, vol. 21, pp. 351-359, 2009.
- [26] A. Zhang, L. Rowland, E. Kaminsky, J. Kretchmer, R. Beaupre, J. Garrett, J. Tucker, B. Edward, J. Foppes, and A. Allen, "Microwave power SiC MESFETs and GaN HEMTs," *Solid-State Electronics*, vol. 47, pp. 821-826, 2003.
- [27] M. Golio, R. Davies, B. Gogoi, D. Lutz, B. Battaglia, R. Neeley, W. Wright, D. Rice, P. Le, and M. Purchine, "A novel high voltage, vertical MOSFET for high power RF applications," in *IEEE Compound Semiconductor Integrated Circuits Symposium, 2008*, 2008, pp. 1-4.
- [28] K. W. Snyder, "An investigation of using a phase-change material to improve the heat transfer in a small electronic module for an airborne radar application," Sandia National Labs., Albuquerque, NM (USA)1990.
- [29] D. Yoo and Y. K. Joshi, "Energy efficient thermal management of electronic components using solid-liquid phase change materials," *IEEE transactions on Device and Materials Reliability*, vol. 4, pp. 641-649, 2004.

- [30] M. Ishizuka, "Operation time control of a high density packaging using a low melting point alloy," *IEEE Transactions on Components and Packaging Technologies*, vol. 27, pp. 239-243, 2004.
- [31] X. Tang, R. Bonner, T. Desai, and A. Fan, "A 2-D numerical study of microscale phase change material thermal storage for GaN transistor thermal management," in *27th Annual IEEE SEMI-THERM Symposium*, 2011, pp. 27-34.
- [32] J. Zeng, Y. Liu, Z. Cao, J. Zhang, Z. Zhang, L. Sun, and F. Xu, "Thermal conductivity enhancement of MWNTs on the PANI/tetradecanol form-stable PCM," *Journal of Thermal Analysis and Calorimetry*, vol. 91, pp. 443-446, 2008.
- [33] J. Zeng, Z. Cao, D. Yang, L. Sun, and L. Zhang, "Thermal conductivity enhancement of Ag nanowires on an organic phase change material," *Journal of Thermal Analysis and Calorimetry*, vol. 101, pp. 385-389, 2010.
- [34] H. Kao, M. Li, X. Lv, and J. Tan, "Preparation and thermal properties of expanded graphite/paraffin/organic montmorillonite composite phase change material," *Journal of Thermal Analysis and Calorimetry*, vol. 107, pp. 299-303, 2012.
- [35] A. Bejan and S. Lorente, *Design with constructal theory*. Hoboken, NJ, USA: John Wiley & Sons, Inc, 2006.
- [36] A. H. Wang, X. G. Liang, and J. X. Ren, "Constructal enhancement of heat conduction with phase change," *International Journal of Thermophysics*, vol. 27, pp. 126-138, 2006.
- [37] J. M. Hill, *One-dimensional Stefan Problems: an Introduction*. New York: Longman Scientific & Technical, 1987.

- [38] V. Alexiades and A. D. Solomon, *Mathematical modeling of melting and freezing processes*. Washington, DC: Hemisphere Publishing Corporation, 1993.
- [39] A. Bejan, *Heat Transfer*. New York: Wiley, 1993.
- [40] O. Bautista, F. Mendez, and I. Campos, "Transient heat conduction in a solid slab using multiple-scale analysis," *Heat and Mass Transfer*, vol. 42, pp. 150-157, 2005.
- [41] B. Gebhart, *Heat Conduction and Mass Diffusion*. New York: McGraw-Hill 1993.
- [42] Y. A. Çengel, *Heat Transfer: A Practical Approach*. New York: McGraw-Hill, 2003.
- [43] "FLUENT 6.3 User's Guide," ed: Fluent Inc., 2006.
- [44] L. Rocha, S. Lorente, and A. Bejan, "Conduction tree networks with loops for cooling a heat generating volume," *International Journal of Heat and Mass Transfer*, vol. 49, pp. 2626-2635, 2006.
- [45] M. B. Healy, H. H. S. Lee, G. H. Loh, and L. Sung Kyu, "Thermal optimization in multi-granularity multi-core floorplanning," in *Design Automation Conference, 2009. ASP-DAC 2009. Asia and South Pacific, 2009*, pp. 43-48.
- [46] *Alloys sorted by temperature*. Available: http://www.indium.com/products/alloy_sorted_by_temperature.pdf
- [47] "International Technology Roadmap for Semiconductos," 2010.
- [48] A. Shakouri and Y. Zhang, "On-chip solid-state cooling for integrated circuits using thin-film microrefrigerators," *IEEE Transactions on Components and Packaging Technologies*, vol. 28, pp. 65-69, 2005.

- [49] P. Wang, A. Bar-Cohen, B. Yang, G. L. Solbrekken, and A. Shakouri, "Analytical modeling of silicon thermoelectric microcooler," *Journal of Applied Physics*, vol. 100, pp. 014501-014513, 2006.
- [50] I. Roch, P. Bidaud, D. Collard, and L. Buchaillot, "Fabrication and characterization of an SU-8 gripper actuated by a shape memory alloy thin film," *Journal of Micromechanics and Microengineering*, vol. 13, p. 330, 2003.
- [51] Y. J. Kim, Y. K. Joshi, A. G. Fedorov, Y. J. Lee, and S. K. Lim, "Thermal characterization of interlayer microfluidic cooling of three-dimensional integrated circuits with nonuniform heat flux," *Journal of Heat Transfer*, vol. 132, p. 041009, 2010.
- [52] K. S. Colinjivadi, J. B. Lee, and R. Draper, "Viable cell handling with high aspect ratio polymer chopstick gripper mounted on a nano precision manipulator," *Microsystem Technologies*, vol. 14, pp. 1627-1633, 2008.
- [53] K. Fukutani and A. Shakouri, "Optimization of thin film microcoolers for hot spot removal in packaged integrated circuit chips," in *22nd Annual IEEE SEMI-THERM Symposium*, 2006, pp. 130-134.
- [54] V. Palankovski, R. Schultheis, and S. Selberherr, "Simulation of power heterojunction bipolar transistors on gallium arsenide," *Electron Devices, IEEE Transactions on*, vol. 48, pp. 1264-1269, 2001.
- [55] G. Humpston and D. M. Jacobson, *Principles of soldering*. Ohio: ASM International, 2004.
- [56] D. Shangguan, *Lead-free solder interconnect reliability*: ASM International, 2005.

- [57] R. Prasher, "Thermal interface materials: historical perspective, status, and future directions," *Proceedings of the IEEE*, vol. 94, pp. 1571-1586, 2006.
- [58] C. E. Green, A. G. Fedorov, and Y. K. Joshi, "Thermal capacitance matching in 3D many-core architectures," in *27th Annual IEEE SEMI-THERM Symposium*, San Jose, 2011, pp. 110-115.
- [59] *eTEC™ HV56 Thermoelectric Cooler Nextreme Thermal Solutions. APS0016 rev 1.0*. Available:
http://nextremethermal.com/media/pdf/techspecs/Nextreme_eTEC_HV56_Data_Sheet.pdf
- [60] J. Stafford, E. Walsh, V. Egan, and R. Grimes, "Flat plate heat transfer with impinging axial fan flows," *International Journal of Heat and Mass Transfer*, vol. 53, pp. 5629-5638, 2010.
- [61] H. Sun, C. Ma, and Y. Chen, "Prandtl number dependence of impingement heat transfer with circular free-surface liquid jets," *International Journal of Heat and Mass Transfer*, vol. 41, pp. 1360-1362, 1998.
- [62] M. Drosig, *Dealing With Uncertainties: A Guide to Error Analysis*. New York: Springer Verlag Berlin Heidelberg, 2007.
- [63] *User's Guide: Dynamic Measurement DC Source Agilent Model 66332A System DC Power Supply Agilent Model 6631B, 6632B, 6633B, 6634B*
Malaysia: Agilent Technologies, 2004.
- [64] *TPS2000B Series Digital Storage Oscilloscopes Service Manual*. Beaverton, OR: Tektronix.

- [65] M. Saini and R. L. Webb, "Heat rejection limits of air cooled plane fin heat sinks for computer cooling," *IEEE Transactions on Components and Packaging Technologies*, vol. 26, pp. 71-79, 2003.
- [66] W. Nakayama, "Exploring the limits of air cooling," *Electronics Cooling*, vol. 12, p. 10, 2006.
- [67] V. Sahu, "Hybrid Solid State/Fluidic Cooling for Thermal Management of Electronic Components," PhD Thesis. Mechanical Engineering, Georgia Institute of Technology, 2011.
- [68] *ShinEtsu Adhesion Promoters*. Available: <http://www.microsi.com/AdhesionPromoters.aspx>
- [69] K. Mackenzie, B. Reelfs, M. DeVre, R. Westerman, and D. Johnson. Optimization of Low Stress PECVD Silicon Nitride for GaAs Manufacturing. *Compound Semi & Microtechnology*, 28-30. Available: <http://www.plasmatherm.com/pdfs/papers/43.%20Optimization%20of%20Low%20Stress%20PECVD%20Silicon%20Nitride.pdf>
- [70] *Wafer Grinding Tapes and Greases*. Available: <http://www.aitechnology.com/products/tapes/grinding-tape/>
- [71] *Application Note: How to Use Fusible Alloys*. Available: <http://www.indium.com/techlibrary/applicationnotes.php>
- [72] S. Guo, I. Lundstrom, and H. Arwin, "Temperature sensitivity and thermal expansion coefficient of benzocyclobutene thin films studied with ellipsometry," *Applied Physics Letters*, vol. 68, pp. 1910-1912, 1996.

- [73] A. Del Campo and C. Greiner, "SU-8: a photoresist for high-aspect-ratio and 3D submicron lithography," *Journal of Micromechanics and Microengineering*, vol. 17, p. R81, 2007.