

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

An Experimental Study of Enhanced Thermal Conductivity Utilizing Columnated Silicon Microevaporators for Convective Boiling Heat Transfer at the Microscale

Permalink

<https://escholarship.org/uc/item/8ks7s202>

Author

Hogue, Christopher William

Publication Date

2011

Peer reviewed|Thesis/dissertation

An Experimental Study of Enhanced Thermal Conductivity
Utilizing Columnated Silicon Microevaporators
for Convective Boiling Heat Transfer at the Microscale

by

Christopher William Hogue

A dissertation submitted in partial satisfaction

of the requirements for the degree of

Doctor of Philosophy

in

Mechanical Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Albert P. Pisano, Chair

Professor Dorian Liepmann

Professor Ming Gu

Fall 2011

Abstract

An Experimental Study of Enhanced Thermal Conductivity Utilizing Columnated Silicon Microevaporators for Convective Boiling Heat Transfer at the Microscale

by

Christopher William Hogue

Doctor of Philosophy in Mechanical Engineering

University of California, Berkeley

Professor Albert P. Pisano, Chair

In this work, silicon-based two-layer microfluidic evaporators were designed, fabricated, and tested in an experimental study designed to examine the effects of out-of-plane vaporization and subsequent lateral convection on the effective in-plane thermal conductivity of the device. The evaporator region in all devices measured 10 mm x 10mm, with an overall device size of 24 mm x 10mm. A columnated structure within the evaporator was utilized to mimic both structurally and functionally the fluid delivery and heat transfer characteristics of the microColumnated Loop Heat Pipe (μ C-LHP) design. In-plane heat flux was provided by a high-power ceramic heater, while working fluid was pumped through the system using a servo-controlled syringe pump.

Temperature data was collected primarily using infrared thermal imaging, and detailed analysis was conducted to estimate convective heat losses from external surfaces and calculate the effective in-plane thermal conductivity. The veracity of this analysis was verified experimentally by analyzing the well-established case of pure conduction and correctly predicting the intrinsic solid conductivity of silicon ($k \sim 130$ W/m \cdot K) over a wide range of input heat fluxes and surface temperatures.

Several distinct performance regimes were observed as the input flux was increased and more vigorous vaporization occurred. Most significantly, subsequent to the onset of stable vaporization, the effective thermal conductivity of a device typically increased by a factor of more than ten ($k \sim 1000$ - 2000 W/m \cdot K). This conductivity was easily maintained and extremely insensitive to further increases in flux. The highest effective thermal conductivities were observed just prior to dryout. Single data point spikes as well as short (3-5 second) intervals of 10,000-20,000 W/m \cdot K were regularly observed.

Peripheral experimental studies related to the author's previous work on self-nucleating surfaces were performed, as was a brief study on the hermetic sealing of microfluidic devices.

Table of Contents

Table of Contents	i
List of Figures and Tables	v
Figures	v
Tables	ix
Acknowledgements	x
Chapter 1: Introduction	1
1.1 Motivation: Planar Thermal Management	1
1.2 The microColumnated Loop Heat Pipe (μ C-LHP): Project Overview	3
1.3 Innovative Features of the μ C-LHP	6
1.3.1 – Coherent Porous Silicon (CPS)	6
1.3.2 – Columnated Vapor Chamber (CVC)	9
1.3.3 – Interline Region Optimized Evaporator Surface (IROES)	12
1.3.4 – Fractal Transport Network (FTN)	13
1.4 Evaporator Optimization	14
Chapter 2: Background on Phase Change Heat Transfer	16
2.1 Elementary Heat Transfer Concepts	16
2.1.1 – Conduction	17
2.1.2 – Convection	17
2.1.3 – Radiation	17
2.2 Phase Change	19
2.3 Nucleation	19
2.4 Boiling Curves	22
2.4.1 – Superheat-Controlled Pool Boiling	23
2.4.2 – Flux-Controlled Pool Boiling	25

Chapter 3: Continued Development of a Self-Nucleating Evaporator Surface	28
3.1 Motivation for Self-Nucleating Evaporator	28
3.2 Self-Nucleating Evaporator (SNE) Concept	31
3.3 Summary of Original Experiments	34
3.4 Modifications to Experimental Setup	38
3.5 New Results	38
3.6 Rationale for Discontinuing SNE Research	41
Chapter 4: μC-LHP Hermeticity Studies	42
4.1 Motivation for Hermetic Sealing	42
4.2 Anodic Bonding	44
4.3 Port Sealing Summary	48
<i>4.3.1 – Solder-based Sealing Technique</i>	<i>49</i>
<i>4.3.2 – Significant Results</i>	<i>52</i>
4.4 Chronological Port Sealing Development	53
<i>4.4.1 – Device Preparation and Filling</i>	<i>54</i>
<i>4.4.2 – Epoxy-based Sealing and Testing Methods</i>	<i>56</i>
<i>4.4.3 – Qualitative Failures and Methodology</i>	<i>57</i>
<i>4.4.4 – Device Layout and Sealant Changes</i>	<i>59</i>
<i>4.4.5 – Quantitative Results</i>	<i>61</i>
Chapter 5: The Open-Loop Columnated Microevaporator	66
5.1 Basic Architecture and Operating Principle	66
5.2 Analogous Heat Transfer and Performance Analysis	70
5.3 Device Layouts and Design of Experiments	78
<i>5.3.1 – Mask Layout Considerations</i>	<i>78</i>
<i>5.3.2 – Floor Designs</i>	<i>82</i>
<i>5.3.3 – Ceiling Designs</i>	<i>84</i>

Chapter 6: Columnated Microevaporator Fabrication	86
6.1 Microevaporator Floor Fabrication	87
6.2 Microevaporator Ceiling Fabrication	95
6.3 Fabrication Minutiae	99
6.3.1 – <i>Through-Wafer Etch Requirements</i>	100
6.3.2 – <i>Unexpected Post-Exposure Bake Effects</i>	101
6.3.3 – <i>Thick Photoresist Mask Exposure Timing</i>	102
6.3.4 – <i>Unintended Backside Etch Stop</i>	104
Chapter 7: Experiments	108
7.1 Experimental Setup	108
7.1.1 – <i>Clamps and Heating System</i>	110
7.1.2 – <i>Infrared Imaging System</i>	111
7.1.3 – <i>Thermistor System</i>	113
7.1.4 – <i>Syringe Pump and Control System</i>	114
7.2 Experimental Strategies and Limitations	116
7.3 Experimental Procedures	118
7.3.1 – <i>Temperature Measurements of IR-Transparent Silicon</i>	118
7.3.2 – <i>Concept Verification: Solid Conduction vs. Vaporization</i>	121
7.3.3 – <i>Primary Conductivity Tests</i>	122
7.4 Heat Transfer Analysis	124
7.5 Notes on Methodology	129
Chapter 8: Results and Discussion	132
8.1 Solid Conduction vs. Vaporization: Concept Verification	133
8.2 Improving the Isothermal Clamp Model	137
8.2.1 – <i>Corrective Temperature Correlations</i>	137

8.2.2 – <i>Effects of Corrective Temperature Correlations</i>	141
8.3 Experimental Verification of Heat Transfer Analysis	149
8.4 General Performance Regimes	150
8.4.1 – <i>Regime I: Dry Baseline</i>	151
8.4.2 – <i>Regime II: Steady Single-Phase</i>	152
8.4.3 – <i>Regime III: Transition</i>	153
8.4.4 – <i>Regime IV: Full Vaporization</i>	155
8.5 Periodicity	156
8.6 Maximum Performance and Design Optimization	158
8.7 Final Comments	160
Bibliography	162
Appendix A: Detailed Discussion of Nucleate Boiling	167
Appendix B: Self-Nucleating Evaporator (SNE) Project	198
Appendix C: Detailed Guide for Columnated Microevaporator Fabrication	239
Appendix D: Sample Calculation of Data Conversion	246
Appendix E: Device Identification Key	252
Appendix F: Heat Transfer Analysis MATLAB Code	254

List of Figures and Tables

Figures

1.1 – Moore’s Law and power density	1
1.2 – Overall Schematic of microColumnated Loop Heat Pipe (μ C-LHP)	4
1.3 – Schematic cross-section of 3-layer device architecture	5
1.4 – Schematic cross-section of device in operation	5
1.5 – Innovative features of the μ C-LHP	6
1.6 – Coherent porous silicon (CPS)	7
1.7 – Traditional wicking materials for heat pipe applications	8
1.8 – Specialized mass flows in the μ C-LHP	11
1.9 – Staged performance degradation of the μ C-LHP	11
1.10 – Regions of the solid-liquid-vapor meniscus	12
1.11 – Example of a fractally branching channel layout	14
2.1 – Electrical circuit analog to multimode heat transfer	18
2.2 – Advancing liquid front and gas entrapment	21
2.3 – The ebullition process	21
2.4 – Superheat-controlled boiling curve during heating	23
2.5 – Superheat-controlled boiling curve during cooling	25
2.6 – Flux-controlled boiling curve during heating	26
2.7 – Flux-controlled boiling curve during cooling	27
3.1 – Vapor burst and backflow	29
3.2 – Nucleation site reactivation sequence without self-nucleation	30
3.2 – Comparison of gas entrapment to self-nucleation	31
3.4 – Molecular structure of azobisisobutyronitrile (AIBN)	32
3.5 – Thermal decomposition of AIBN	32
3.6 – Nucleation site reactivation sequence with self-nucleation	33
3.7 – Expected qualitative results of cyclic boiling experiments	34
3.8 – Photographs and schematic cross-sections of submersible test platform	35
3.9 – Schematic of double boiler apparatus	36

3.10 – Photograph of SNE test setup	36
3.11 – Schematic of modified heater block	38
3.12 – Successful cyclic boiling test using modified experimental setup	40
4.1 – Water permeability through various solids	43
4.2 – Anodic bonding physics	45
4.3 – Characteristic current behavior in anodic bonding	45
4.4 – Anodically bonded wafers with voids	46
4.5 – Mitigation of voids using increased bond width	46
4.6 – Karl Suss SB6 Thermocompression and Anodic Bonder	47
4.7 – Schematic of metalized fill-port fabrication process flow	49
4.8 – Metalized fill ports	50
4.9 – Solder-based sealing technique	51
4.10 – Solder-sealed devices	52
4.11 – Exploded solder-sealed devices	53
4.12 – Schematic and photograph of Phase I μ C-LHP	54
4.13 – Esec 8003 Dicing Saw	55
4.14 – Epoxy-based sealing techniques	56
4.15 – TPS Tenney Junior Environmental Test Chamber	57
4.16 – Microcrack propagation in Pyrex	58
4.17 – HF-etched Pyrex wafer	58
4.18 – Hermeticity-specific device and wafer layouts	60
4.19 – Filled devices sealed with E6000 epoxy	61
4.20 – Successful 90-day hermeticity test results	62
4.21 – Failed 57-day hermeticity test results	63
4.22 – 50x micrographs of well- and poorly-sealed ports	64
5.1 – Exploded view and cross-section of open-loop columnated evaporator	66
5.2 – Dimensioned drawings of silicon components	67
5.3 – Cross-sectional schematic of microevaporator operation	68
5.4 – Theoretical motivation of effective conductivity measurements	69
5.5 – Comparison of μ C-LHP to the open-loop microevaporator	71
5.6 – Geometry and boundary conditions of simplified analytical model	72

5.7 – Theoretical temperature profiles as a function of convection coefficient	75
5.8 – Theoretical temperature gradients as a function of convection coefficient	76
5.9 - Theoretical thermal conductivity as a function of convection coefficient	76
5.10 – Schematic of the effect of columnation scheme on intercolumnar film thickness	77
5.11 – Wafer-level mask layout for device floor fabrication	80
5.12 – Wafer-level mask layout for device ceiling fabrication	81
5.13 – Columnation parameter definitions	82
5.14 – Schematic of evaporator floor columnation schemes	83
5.15 – Geometry and channel parameters of ceiling microtextures	85
6.1 – Device floor/ceiling nomenclature	86
6.2 – Schematic wafer cross-sections for floor process flow (Steps 1-4)	87
6.3 – Characteristic dark-field photomask at device level	88
6.4 – Schematic wafer cross-sections for floor process flow (Steps 5-8)	89
6.5 – Micrographs of floor features after hard-baking	90
6.6 – Schematic of wafer-floating for backside oxide etch	90
6.7 – Two causes of etch-rate non-uniformity	91
6.8 – Schematic wafer cross-sections for floor process flow (Steps 9-11)	91
6.9 – Partially etched 400 μm columnar holes	92
6.10 – Micrographs of floor features after thick resist removal	92
6.11 – Schematic wafer cross-sections for floor process flow (Steps 12-15)	93
6.12 – Nearly completed 6” floor wafer	94
6.13 – Diced floor devices	94
6.14 – Schematic wafer cross-sections for ceiling process flow	95
6.15 – Photomask for patterning ceiling microtextures	96
6.16 – Micrographs of all eight parallel channel ceiling microtextures	97
6.17 – Micrographs of all eight orthogonal channel ceiling microtextures	98
6.18 – Micrographs of all eight oblique channel ceiling microtextures	99
6.19 – Diced floor and ceiling components	99
6.20 – Anisotropic etching of silicon using KOH	100
6.21 – Cracks in thick photoresist	101
6.22 – Ring-like appearance in thick photoresist windows	102

6.23 – Diminished ring effect	102
6.24 – Severely underexposed thick photoresist	103
6.25 – Slightly underexposed thick photoresist	103
6.26 – Properly exposed thick photoresist	103
6.27 – Schematic wafer cross-sections for earlier floor process flow (Steps 10-14)	105
6.28 – Micrograph of successful punch-through	106
6.29 – Steep etch gradients and grassy floor deposits	107
6.30 – Broken column sidewalls	107
7.1 – Conceptual schematic of experimental setup	108
7.2 – Photograph of experimental setup	109
7.3 – Enlarged view of two-layer device and clamps	109
7.4 – Schematics and photographs of device clamps	110
7.5 – Heater clamp design with insulating air gap	111
7.6 – Screenshot of ExaminIR software and profiles plots	112
7.7 – Screenshot of ExaminIR software and statistical tools	112
7.8 – Data acquisition system and thermistor bridge circuitry	114
7.9 – Screenshot of Labview virtual instrument	115
7.10 – Heater tab overlap variation	117
7.11 – Sequential IR images demonstrating transparency of silicon	119
7.12 – Use of Kapton tape for IR temperature measurement	120
7.13 – Effect of Kapton tape resistance on temperature measurement	120
7.14 – Statistical domains monitored during concept verification tests	121
7.15 – ExaminIR screenshot with linear and areal domains for conductivity tests	126
7.16 – Flowchart illustrating the heat transfer analysis for conductivity calculation	128
8.1 – Variation in domain temperatures in a representative dry test	134
8.2 – Temperature gradient and standard deviation in a representative dry test	134
8.3 – Variation in domain temperatures in a representative wet test	136
8.4 – Temperature gradient and standard deviation in a representative wet test	136
8.5 – Multi-view IR images for clamp surface temperature correlation	139
8.6 – Multi-view temperature data plotted to correlate clamp surface temperatures	140
8.7 – Convective losses: isothermal approximation vs. correlated surface temperatures	144

8.8 – IR images and color maps for central/peripheral partition	145
8.9 – Graphical view illustrating a range of central/peripheral partitions	146
8.10 – Effect of weighting factors at low power	148
8.11 – Effect of weighting factors at high power	148
8.12 – Simplified polynomial convective loss correlations	149
8.13 – Variation in surface temperature during dry testing	150
8.14 – Corresponding variation in calculated effective thermal conductivity	150
8.15 – Temporal plot of thermal conductivity in Regime I	151
8.16 – Temporal plot of thermal conductivity in Regime II	153
8.17 – Temporal plot of thermal conductivity in Regime III	154
8.18 – Temporal plot of thermal conductivity in Regime IV	155
8.19 – Temporal plot of thermal conductivity for full heating cycle	156
8.20 – Representative temporal plot illustrating approximate periodicity	157
8.21 – Additional test examples illustrating periodicity	157

Tables

1.1 – DARPA Thermal Ground Plane program goals	3
3.1 – Summary of boiling parameters during cyclic boiling tests	37
4.1 – DARPA Thermal Ground Plane program hermeticity milestones	43
5.1 – Summary of floor columnation schemes	84
8.1 – Heater clamp: isothermal approximation vs. correlated surface temperatures	142
8.2 – Device clamp: isothermal approximation vs. correlated surface temperatures	142
8.3 – Convective losses: isothermal vs. correlated temperatures at low power	147
8.4 – Convective losses: isothermal vs. correlated temperatures at high power	147

Acknowledgments

I completed my M.S. Thesis at the end of a very turbulent and seemingly impossibly influential period of my life. I had never before experienced such a roller coaster of emotions nor seen such dramatic changes in my own personality and world view. Consequently, my acknowledgments digressed into a longwinded, flowery essay that drifted from religious faith to philosophy to sociopolitical theory. Things are very much different now, and I am finally closing one exceptionally long chapter in my life. After twelve (!) years at Berkeley—twelve of the most fantastic, magical years I can possibly imagine—this is my final text, and for once I shall strive for brevity.

Thank you to the University, thank you to the College, thank you to the City of Berkeley and the City by the Bay. Thank you to the hills and the pathways and my beloved fog.

Thank you to my friends and colleagues, my instructors and mentors.

Thank you especially to Professor Albert Pisano. Words cannot express my gratitude.

To my family, your love and support is unfathomable.

And to the LORD, I can only tremble in awe.

This research was made possible primarily through funding from the Defense Advanced Research Projects Agency (DARPA), in association with the Microsystems Technology Office (MTO) and the Thermal Ground Plane (TGP) Program, issued by DARPA/CMO under Contract No: HR0011-10-C-0100.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Chapter 1: Introduction

This chapter is divided into three major sections. In the first, I discuss the problem of thermal management in modern high-power electronics, as well as the general motivation for this research. In the second section, I describe the specific approach taken by our research team to address this problem and introduce the overall architecture and operating principles of our novel substrate, the microColumnated Loop Heat Pipe (μ C-LHP). Finally, in the third section, I discuss the elements of this multifaceted project for which my own personal research is geared and on which this study is primarily focused, namely evaporator optimization.

1.1 Motivation: Planar Thermal Management

Since the late 1960s, Moore's Law, which states that the transistor density of integrated circuits will double approximately every two years, has proven remarkable accurate (see Fig. 1.1a). The great reliability of this trend led computer designers toward a paradigm whereby performance was maximized by utilizing a single monolithic processor with the highest transistor density, the largest internal cache, and the fastest clock speed available. An inescapable consequence of the ohmic heating associated with denser transistor arrays has been a dramatic increase in the thermal energy dissipated by these chips (Fig. 1.1b). Quite naturally, the great increases in *areal* power density have demanded more effective cooling systems to prevent thermal failure. However, because denser, faster chips kept rolling off the fabrication line with such regularity, it didn't make sense to keep redesigning more efficient cooling systems to dissipate the increased thermal loads. Instead, designers simply modified old system-level architecture to keep the *volumetric* power density more or less constant. The archetypal example is decreasing the number of blades per volume in a server bank (and/or increasing the airflow).

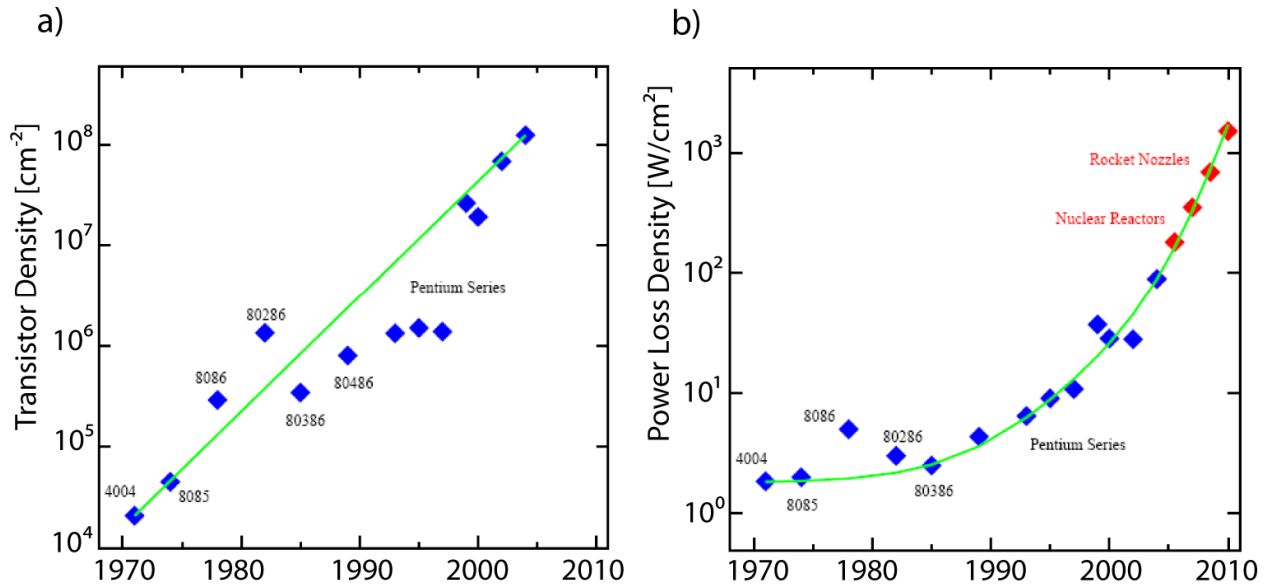


Fig. 1.1 – a) Historical verification of Moore's Law and **b)** associated power densities (Source: ITRS).

Until very recently, this paradigm has satisfied the demand for increased computing power, largely because Moore's Law has proven so reliable. However, some chip designers feel they have reached the point of diminishing returns when it comes to transistor density. Leakage current from inactive transistors is now proving to be the limiting factor for further increases in transistor density—likely spelling the end for Moore's Law. Moreover, if the historical rate of increase in transistor density slows, so too does the associated rise in areal power density. This means that there is now a much greater incentive on the part of systems designers to reconsider the old thermal paradigm. Specifically, if the areal power density of individual blades no longer increases so dramatically from year to year, perhaps it makes more sense to design systems that utilize *multiple moderately-powered processors that occupy the same thermal envelope as a single state-of-the-art monolithic processor*. In other words (and again invoking the server bank example), rather than decreasing the number of blades per volume, the new paradigm would instead increase the blade count, but make each blade more efficient at dissipating heat.

This then brings us to the real issue: *substrate-level heat dissipation*. Metallic and even synthetic diamond substrates, which rely solely on solid conduction for thermal transport, are gradually proving themselves inadequate as heat spreaders for closely spaced substrates. Single-phase liquid-cooled systems offer much improved performance over solid conduction alone, but they require a large volume of working fluid to be effective, as well as bulky power and control systems. Consequently, there is tremendous interest in cooling technologies that utilize liquid-to-vapor phase change for the absorption and rejection of waste heat. Not only does phase change dramatically increase effective thermal conductivity by taking advantage of convected latent heat (with much less working fluid), it also allows for near isothermal heat transfer, thus promoting homogeneous thermal conditions (*i.e.* low ΔT) across large, often thermally heterogeneous chip carriers. Capillary-driven systems such as heat pipes and thermosyphons have garnered particular interest because of their simple design and long history of reliability. Unlike vapor compression systems, heat pipes contain no moving parts, require no external power supply or interconnect, and operate completely passively. As the name implies, however, traditional heat pipes are indeed long cylindrical tubes that transport heat from one end to the other—a geometry quite ill-suited to the highly planar structure of electronic substrates.

Recently, however, companies have developed so-called planar heat pipes, essentially flat, hollow evaporation chambers with a porous internal coating, to more evenly distribute thermal energy beneath traditional finned heat sinks. Unfortunately, these systems only challenge the limits of thermal performance by straightforward “macroscale miniaturization” of decades-old heat pipe technology. They have thus far failed to utilize the broad range of photolithography-based fabrication techniques and precision surface control available at the microscale. In addition, the thermal expansion mismatch of their metallic shells makes direct, ultra low-resistance thermal bonding to semiconductor materials nearly impossible. This is especially true for systems which encounter continuous thermal cycling over large temperature ranges. In summary, though commercially-available planar vapor chambers demonstrate commendable improvement over traditional solid metallic heat spreaders and liquid-cooled systems, they do not offer fundamentally groundbreaking or adaptable technology to address the high-performance needs of the future. The Microsystems Technology Office (MTO) of the Defense Advanced Research Projects Agency (DARPA) also recognized this need and consequently issued a Broad Agency Announcement (BAA) in the Spring of 2007 calling for novel technological proposals for high-performance heat spreading substrates.

This program is known as the Thermal Ground Plane (TGP), a reference to the completely analogous *potential* ground plane common to many electronic circuits. In much the same way that a potential ground plane provides a uniform “sink” fixed at 0 V, a substrate acting as a thermal ground plane would have such high in-plane thermal conductivity (*i.e.* heat spreading ability), that the entire surface would be fixed at a uniform temperature, namely the saturation temperature of the working fluid. In terms of definable metrics, the primary goals of the TGP program are summarized below in Table 1.1.

Table 1.1 – DARPA Thermal Ground Plane Program Goals

Thermal Transport	Effective thermal conductivity: $k \sim 20,000$ W/m·K (compared to ~ 200 W/m·K for state-of-the-art copper-molybdenum metallic substrates)
Operation Under Acceleration	Operate at target performance under 20g loading
Planar Geometry	Thickness < 1 mm; Lateral dimensions > 10 cm
Low Density	< 3000 kg/m ³ ($\sim 1/3$ that of copper)
Thermal Expansion Match	Matched within 1% to Si (2.6 ppm/K), GaN (3.2 ppm/K), or GaAs (6.9 ppm/K)
Hermeticity	< 0.1 % fluid loss per year (at 100°C)
Operation Duration	1000 hours

1.2 The microColumnated Loop Heat Pipe (μ C-LHP): Project Overview

In response to this BAA, a team of researchers from the University of California, Berkeley (including the present author), put forth a proposal we have called the microColumnated Loop Heat Pipe, or μ C-LHP. The μ C-LHP concept, shown schematically in Fig. 1.2 below, combines the proven simplicity and reliability of capillary pumped loop systems, with the design, innovation, and integration expertise of a leading university MEMS research laboratory. Rather than simply continuing the trend of “top-down” miniaturization, this design is based on the principle of “bottom-up” fabrication, fully utilizing recent advances in microtechnology.

The entire device employs a bonded three-layer architecture consisting of an Evaporator layer, a Wick and Vapor Transport layer, and a Liquid Transport layer (see Fig. 1.3). The Evaporator layer serves dual purposes by 1) providing a micro-patterned, liquid-spreading surface for maximized interline evaporation and 2) capping and sealing the columnated vapor chamber. Below the Evaporator layer is the Wick and Vapor Transport layer, the most novel and adjustable element of the μ C-LHP. The central region of the layer (*i.e.* the evaporator footprint) will be coherent porous silicon (see Section 1.3.1), with the vapor space etched away entirely, leaving an array of porous silicon columns rising from a porous silicon floor. Leading away from the evaporator, a series of channels branch and divide in a fractal-like manner, transporting vapor to the condenser. Finally, at the bottom of the stack, the Liquid Transport layer consists of a central reservoir beneath the evaporator footprint and several return lines to transport liquid back to the reservoir.




Fig. 1.2 – Overall schematic of the UC-Berkeley-proposed microColumnated Loop Heat Pipe system, or μ C-LHP.

The μ C-LHP operates as follows (see Fig. 1.4):

- 1) Capillary pressure pulls liquid up from the reservoir through the Primary Wick.
- 2) Some liquid evaporates as it exits the floor pores and sidewalls of the tapered columns, cooling the Primary Wick and keeping the vapor chamber saturated.
- 3) The rest of the liquid continues to travel all the way up the Columnar Wick to spread over the evaporator surface, cooling the chip as the thin film vaporizes.
- 4) Vapor exits the vapor chamber via vapor outlet channels, where it travels away from the evaporator toward the condenser via the fractally-branching transport lines.
- 5) After condensing, the liquid is drawn back to the reservoir via wick capillary pressure to continue the cycle.

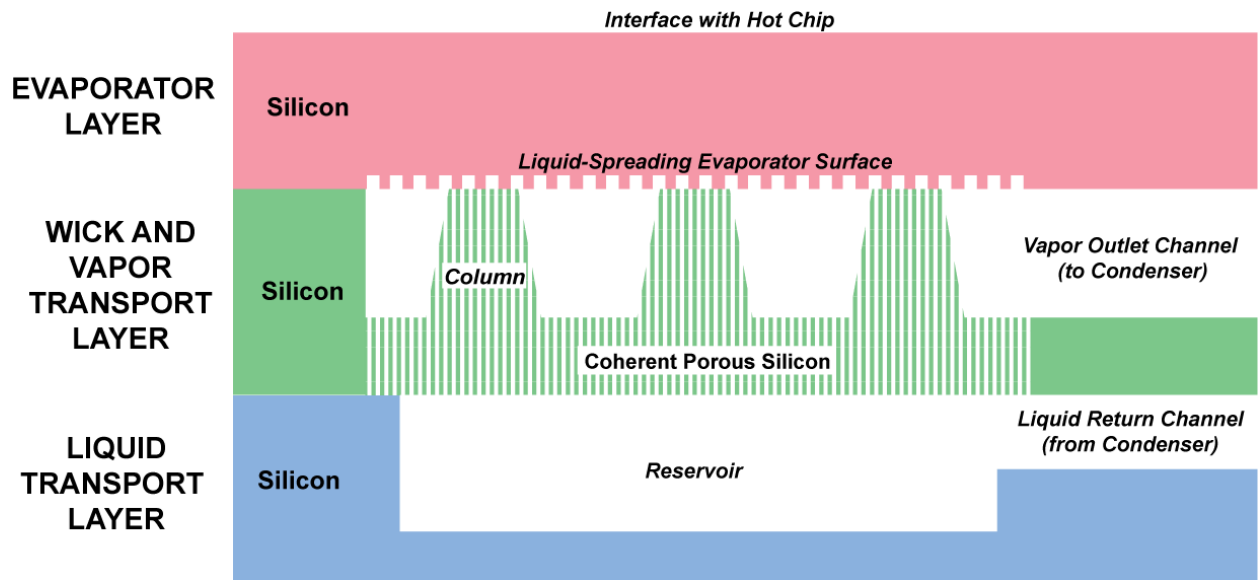


Fig. 1.3 – Schematic of three-layer architecture, highlighting the evaporator cross-section.

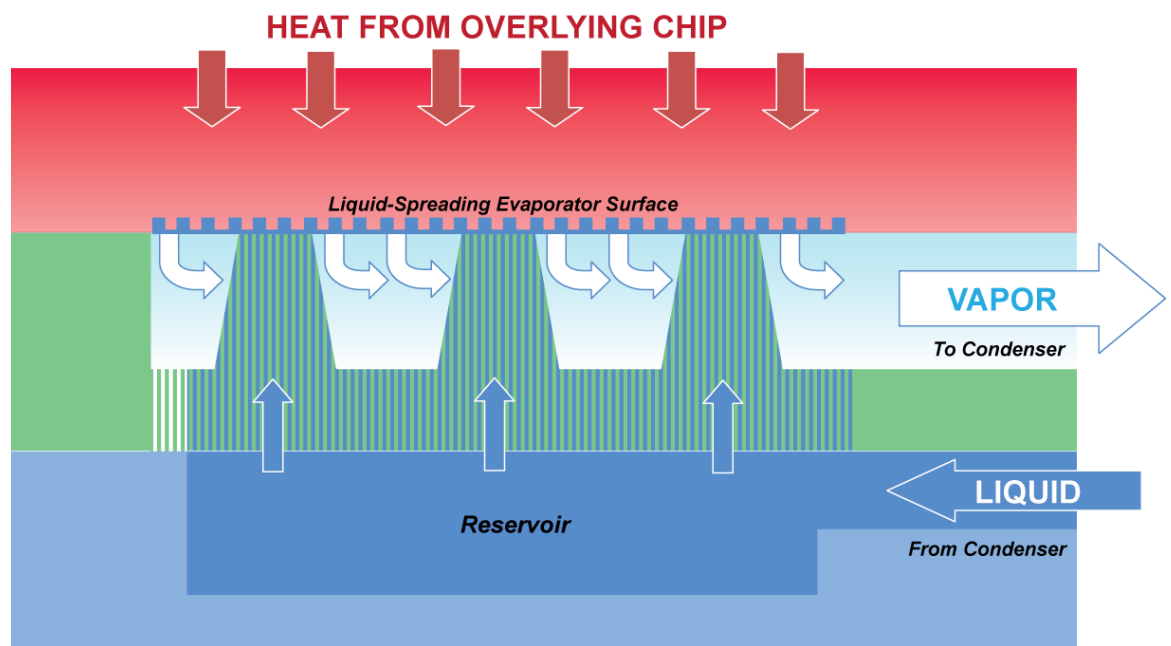


Fig. 1.4 – Schematic of evaporator cross-section during operation.

1.3 Innovative Features of the μ C-LHP

There are four primary innovative features of the μ C-LHP, highlighted in Fig. 1.5 below.

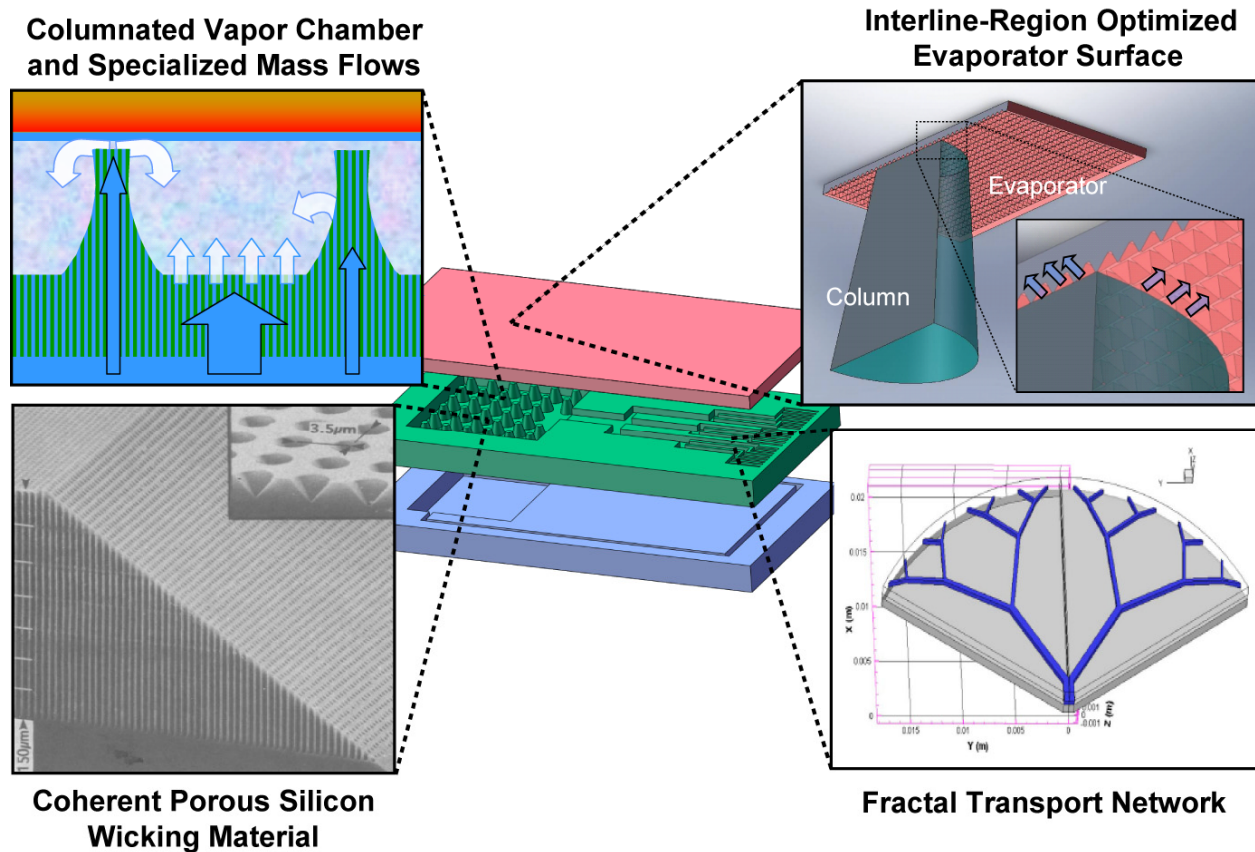


Fig. 1.5 – The four primary innovative feature of the μ C-LHP design.

1.3.1 – Coherent Porous Silicon (CPS)

First and foremost is the unique wicking material: coherent porous silicon, often abbreviated simply as CPS. CPS consists of bulk single-crystal n-type silicon whose surface has been electro-chemically etched in hydrofluoric acid, typically utilizing backside illumination to lower the required electric field strength needed to generate excess holes (the minority charge carrier in n-type silicon) at the pore tip. The holes concentrated at the pore tip lead to further dissolution of the silicon at the tip, while the n-type doping of the bulk substrate gives rise to a depletion of holes along the sidewalls, thus passivating them against further dissolution in the electrolyte. Pores orthogonal to the surface, with depths up to the wafer thickness, diameters as large as 10 μm , and with aspect ratios as large as 250 have been produced by this technique. (Please note that this is only a cursory description of a highly complex phenomenon. For a more detailed discussion of macropore formation in n-type silicon, please see Lehmann's excellent quantitative study [1.1]). Extensive mathematical modeling and experimental verification of this phenomenon have been performed by Lehmann and others, and straightforward equations enable

us to accurately construct the desired pore geometry based on a few formation parameters such as doping density, applied bias, illumination intensity, and HF concentration.

Without pre-processing of the bulk silicon, the electrochemical etching process will produce a random pattern of varying-sized macropores, ultimately dependent on the intrinsic surface roughness of the silicon and the applied current density (see Fig. 1.6a). However, by first pre-patterning the silicon surface using standard photolithographic masking and etching, this random pattern can be made regular and tightly controlled, both in pore spacing and diameter (Fig. 1.6b). Specific areas can be completely masked off from the electrochemical etching, and moreover, the bulk CPS itself can later be etched to create three-dimensional structures made entirely of CPS. Although it was the surprising optical properties of CPS that first garnered the attention of the scientific community in the early 1990s [1.2, 1.3], others (the present author included) took primary interest in the high aspect-ratio pores as a possible wicking material for capillary-driven mass transport, particularly researchers at the University of Cincinnati [1.4, 1.5].

The “parametric” aspect of CPS, however, is perhaps its greatest asset to the design process, particularly in the optimization stages. Accurate control over pore size and density will allow us to precisely adjust capillary force and flow resistance within the wick, providing us with parametric “knobs” to fine-tune the overall wick performance. Finally, it must be pointed out that the mass transport properties of CPS have been characterized and verified experimentally by Lehmann and others. The work of Hoelke, *et al.* [5] has proven especially valuable for modeling the relationship between pore size and capillarity, as well as developing mathematical expressions for wick permeability. Their models have been advanced further by Hamdan *et al.* [1.6, 1.7] and Dhillon [1.8], who have added parametric studies of wick parameters to better determine overall trends in loop heat pipe performance.

Not only is the pore geometry highly-controllable, but the native silicon dioxide present on the walls of the pores provides a natural hydrophilic surface conducive to capillary-driven flow. While both traditional macroscale heat pipes and planar miniaturized heat pipes have relied almost exclusively on bulk-machined or “stochastic” wicking materials such as sintered metallic powders, fabrics, or foams (Fig. 1.7), the near-perfect regularity of coherent porous silicon allows for much more controllable design and fabrication, as well as simpler, more characteristic modeling of capillary-driven phenomena. In addition, unlike sintered metal powders or foams (which would require entirely new microfabrication processes), the formation of CPS from standard n-doped wafers is easily incorporated into our fabrication process. In fact, the porous nature of CPS actually makes definition and etching of the vapor space and columnated wick structure even easier. Finally, by utilizing a silicon-based substrate material (rather than metallic), we naturally avoid any thermal expansion mismatch with directly bonded electronics or other system components. These advantages do not come without tradeoffs. The bottom-up, microfabrication-intensive strategy of the μ C-LHP is undoubtedly more complex and expensive from a *manufacturing* standpoint. However, we firmly believe that the tremendous gains in simplicity and performance optimization from a *design* standpoint will ultimately outweigh the costs.




Fig. 1.6 – a) Coherent porous silicon (CPS) etching on a 45° bevel without pre-processing; **b)** CPS etching on a 45° bevel utilizing a regular pattern of pre-etched starter holes (from [1.1]).




Fig. 1.7 – Two typical examples of traditional wicking materials used in most macroscale and miniature heat pipe applications: a) Sintered metallic powder; **b)** Sintered metallic fabric.

1.3.2 – Columnated Vapor Chamber (CVC)

Although coherent porous silicon, in and of itself, offers tremendous potential as a wicking material for heat pipe applications, what truly differentiates the μ C-LHP from other CPS-based heat pipes systems is the distinctive columnated structure of the evaporator. As can be seen in Fig. 1.2, the evaporator region of the device looks much like the inside of a cave, with floor-to-ceiling columns resembling joined stalactites and stalagmites. Both the floor and the columns themselves are made entirely of CPS (with pores oriented vertically and etched through the entire wafer), while the empty space between them has been subsequently removed using a standard silicon etching process. Thus there is a solid continuous wick from the underlying reservoir to the floor of the chamber, with a fraction of the wick extending all the way to the ceiling via the columns. Due to the tapering of the columns, some of the pores near the periphery of the columns will open to the chamber along the column sidewalls, while the majority of the pores in the core region of the columns carry liquid all the way to the ceiling. This Columnated Vapor Chamber (CVC), as we call it, creates a truly three-dimensional wicking structure that provides several distinct advantages.

The most critical aspect of continuous heat pipe operation is ensuring wick saturation. In non-columnated designs, the high-temperature evaporator surface is in direct contact with the primary wick. If the heat flux temporarily increases above the normal operating range, conduction through the wick may lead to elevated wick temperatures sufficient to cause complete dryout of the wick. This effectively terminates capillary pumping and leads to device failure. The CVC design, on the other hand, separates evaporator dryout from wick dryout by disengaging the Evaporator from the Primary Wick. Only a designated fraction of the wicked liquid is transported through the Columnar Wick to the high-temperature Evaporator surface (see Fig. 1.8a). The remainder evaporates from the Primary Wick “floor” (under much more controlled conditions) directly to the vapor space (see Fig. 1.8b). Any transient increases in heat flux from the Evaporator will not dramatically alter the bulk vapor conditions in the chamber, thus attenuating any changes in the evaporation rate from the Primary Wick. In addition, evaporative cooling from pores along the sidewalls of the tapered columns provides an alternate path for conduction down the columns, thus redirecting heat away from the Primary Wick (see Fig. 1.8c). In summary, the CVC architecture provides three specialized mass flows to ensure more robust operation. One mode transports liquid up the columns to cool the actual Evaporator, the second keeps the Primary Wick and vapor space saturated, while the third prevents potentially dangerous conduction down the columns to the Primary Wick.

By ensuring the saturation of the Primary Wick, the CVC provides the μ C-LHP with unique performance characteristics. During normal operating conditions, capillary-driven liquid transport through the Columnar Wick is sufficient to replenish the thin-film on the Evaporator surface. We call this ‘Standard Flux Mode’ (see Fig. 1.9a). If the applied heat flux increases, however, the rate of evaporation will eventually exceed the rate at which the capillary forces can supply fluid to the Evaporator. As a result, the thin liquid film on the Evaporator surface will begin to dry out and the Evaporator wall temperature (T_w) will increase. With the CVC’s columnated structure, however, the Primary Wick can remain saturated even as the Evaporator surface dries out, ensuring that system-wide capillary pumping continues. While the Evaporator must transition to a higher operating temperature to compensate for surface dryout, the attenuation of Primary Wick evaporation provided by the saturated vapor space provides a

second stable, albeit temporary, “Hi-Flux Mode” (see Fig. 1.9b). In this mode, sufficient response time would be available to power down the electronics and avoid thermal damage. On the other hand, with traditional non-columnated wicking structures, such an increase in the applied flux leads to unmitigated conduction through the primary wick. As a result, the wick dries out entirely, system-wide capillary pumping ceases, and an exponential increase in wall temperature leads to thermal failure of the overlying electronics (see Fig. 1.9c). Essentially, the columnar wicking structure of the CVC separates *evaporator* dryout from *wick* dryout, allowing for “staged” performance degradation, rather than catastrophic failure, in extreme operating conditions.

Because of these specialized mass flows discussed above, the Columnated Vapor Chamber provides yet another huge advantage for parametric design, troubleshooting, and optimization. Just as the highly-controllable CPS structure allows us precise control over the size and density of the pores for optimized wicking, the CVC architecture allows us precise control over the column diameter, column height, and inter-columnar spacing. This provides us with yet another series of “knobs” to fine-tune the overall wick performance. For example, increasing the ratio of column-to-floor “footprint” (by increasing the diameter or density of the columns) will provide more cooling to the Evaporator, while decreasing this ratio will provide more cooling to the Primary Wick. This allows us great flexibility for properly distributing mass flow to counter the observed failure mode (*e.g.* Evaporator dryout or Primary Wick dryout).

Another common problem with traditional non-columnated loop heat pipes is the long, narrow channels through which vapor must escape to reach the vapor transport lines. In the micromachined heat pipe proposed by Hoelke *et al.* [1.5], a series of parallel “vapor-removing ducts” were added to the bottom surface of the capping wafer to provide exit paths for the vapor (since the solid CPS wick abutted directly against the capping wafer). In their model, as the liquid evaporated from the wick, the tremendous volumetric expansion led to extremely high vapor velocities in the ducts, justifying turbulent flow regimes. Because of the solid wick, the nominal path length for the vapor ducts was approximately half the length of the evaporator, causing huge pressure drops that dominated the total system losses. A columnated wick, on the other hand, largely eliminates this problem. While the very fine pores *within* the CPS columns will still generate the large capillary forces necessary for passive pumping, the large open “vapor space” *between* the columns allows the newly vaporized working fluid to travel relatively unimpeded through the vapor chamber to the outlet transport channels. This reduces the nominal path length for the escaping vapor to the radius of the columns and prevents the large pressure drop associated with high-velocity vapor being forced through highly-constrictive porous wick structures. This ultimately allows for higher mass flow rates and thus greater heat dissipation.




Fig. 1.8 – Three specialized mass flows provided by the unique columnated structure of the Columnated Vapor Chamber (CVC) allow for more robust operation.

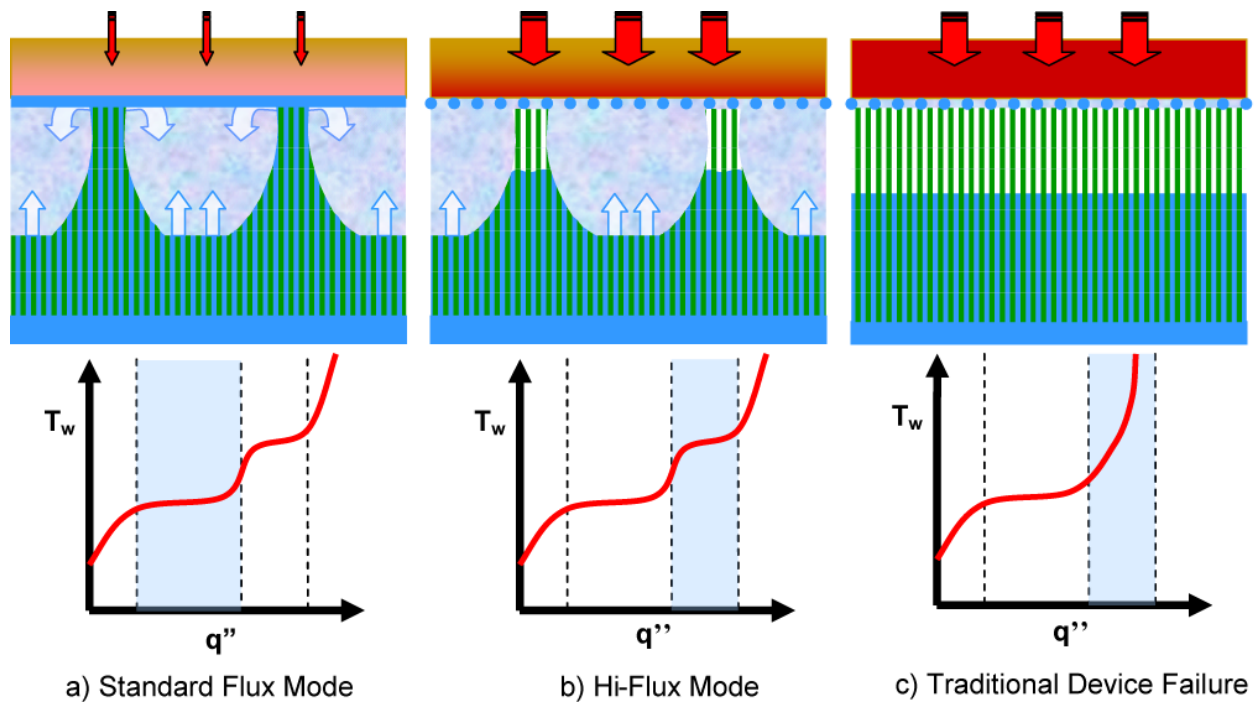


Fig. 1.9 – Staged performance degradation in the μ C-LHP (a and b), versus complete wick dryout and catastrophic failure in non-columnated systems (c).

1.3.3 – Interline Region Optimized Evaporator Surface (IROES)

The third innovative feature of the μ C-LHP, the Interline Region Optimized Evaporator Surface (or IROES for short), is an evaporator surface designed to interface with the columnar wick in such a way as to maximize the total surface area of what is known as the *interline region*. At a solid-liquid-vapor interface, illustrated schematically in Fig. 1.10, the meniscus can be divided into two major regions. In the “macro” region (also called the *intrinsic meniscus*), the curvature of the interface is determined almost entirely by capillary forces, and the liquid film is sufficiently thick to greatly inhibit heat transfer via conduction across the film. This large thermal resistance means that very little evaporation (and thus heat transfer) will occur in the macro region.

Moving away from the intrinsic meniscus, the “micro” region (also called the *extended meniscus*) can be divided into two regions, the *adsorbed region* and the *interline region*. Throughout the extended meniscus, the film is sufficiently thin that attractive dispersion forces (sometimes called *Lifshitz-van der Waals forces*) between the solid and liquid and the liquid molecules themselves contribute significantly to the curvature. At the extreme end, in the adsorbed region, these forces are so strong that they completely prohibit evaporation from the surface, regardless of the wall temperature. However, between this adsorbed region and the intrinsic meniscus lies the interline region, where the film is thin enough that the conductive thermal resistance is vanishingly small, but the dispersion forces are not yet strong enough to bind the liquid molecules to the surface. Consequently, heat and mass transfer coefficients in the interline region are up to three orders of magnitude higher than their average values over the entire meniscus [1.9], and up to half of the total evaporation and accompanying heat transfer occur in this small region [1.10, 1.11, 1.12].

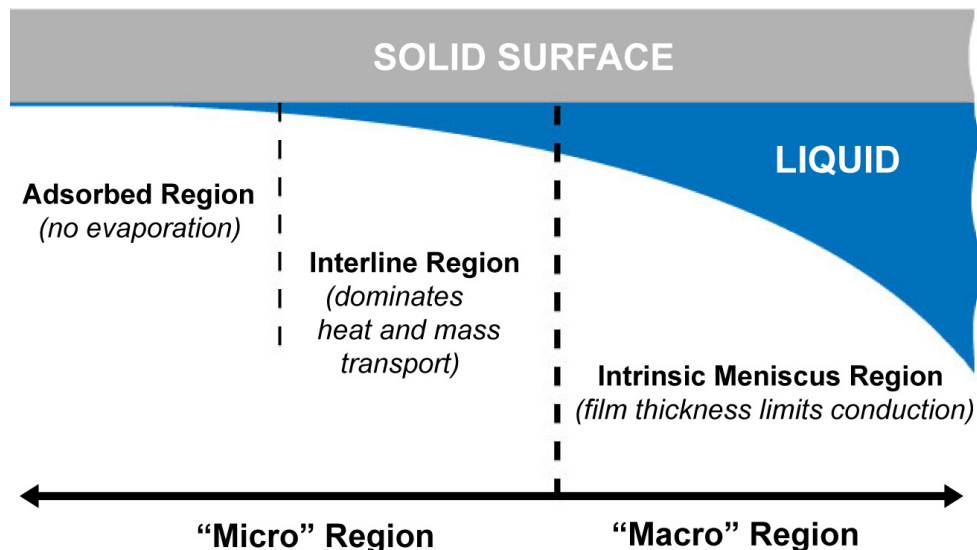


Fig. 1.10 – Schematic illustrating the various regions making up a typical solid-liquid-vapor interface. (Note: purely for illustrative purposes; regions are not drawn to scale.)

1.3.4 – Fractal Transport Network (FTN)

The fourth and final innovative feature of the μ C-LHP is known as the Fractal Transport Network, or FTN. In simplest terms, it consists of a branching network of progressively smaller-diameter microchannels designed to transport the vaporized working fluid from the evaporator to the condenser, while minimizing the pressure drop. As discussed earlier in Section 1.3.2, one of the major limitations of microscale loop heat pipes is the large pressure drop associated with high-velocity vapor flow through small-diameter channels. Although turbulent flow is almost never encountered in single-phase microfluidic systems, the enormous volumetric expansion associated with liquid-to-vapor phase change can easily lead to meter-per-second flow rates in the vapor lines, along with the massive frictional losses that accompany turbulence. Consequently this single term tends to dominate the total pressure drop throughout the entire closed loop. Since this pressure drop must ultimately be offset by the positive capillary pressure provided by the wick, it is often the limiting factor on the maximum mass flow rate and thus the maximum heat dissipation.

From a *mass transfer* perspective, then, it is usually desirable to have large-diameter vapor transport lines to decrease the vapor velocity and thus decrease the pressure drop. On the other hand, from a *heat transfer* perspective, it is usually desirable to have small-diameter channels to maximize the surface-to-volume ratio and more efficiently spread the convected thermal energy over a larger surface area. These competing aims have led researchers to develop biologically-inspired branching patterns (see Fig. 1.11) that minimize pressure drop for a given input heat flux and spreading area. This aspect of the μ C-LHP project is probably least significant to my research, however, so I will only briefly mention a few examples.

Using a one-dimensional model, Pence [1.13] predicted the pressure distribution and wall surface temperature along a fractal-like branching network, compared with an array of parallel channels having identical flow length, pressure drop, and pumping power. Results indicate that the fractal network requires only 50% of the convective area of the parallel network for comparable performance. This one-dimensional model was validated by Alharbi *et al.* [1.14] using three-dimensional CFD analysis and experiments. More recently, Wang *et al.* [1.15] compared fractal-like networks to parallel networks using the CFD package, Fluent 6.2, to study additional effects such as channel blockage. Not only did the results indicate a near 50% decrease in pressure drop and significantly higher average Nusselt numbers for the fractal network, they also found the fractal network to be much less sensitive to channel blockage. Network-level optimization schemes have been proposed by Wechsato *et al.* [1.16] to determine optimal branching angles, segment lengths, and channel diameter ratios to minimize pressure drop for a given number of total branching levels and branching pattern. For example, it has been shown that global flow resistance will be minimized if the branching pattern adheres to the following simple algorithms:

$$\frac{d_{k+1}}{d_k} = n^{-1/3} \qquad \frac{L_{k+1}}{L_k} = n^{-1/2}$$

where d and L are the hydraulic diameter and length of a particular channel segment, n is the number of “daughter” branches at each split, and k and $k+1$ represent the lower and higher order branching level, respectively. In the final iterations of this project, we will adopt these or similar

algorithms to help minimize the pressure drop in the vapor transport lines and condenser to increase the maximum allowable mass flow rate for a given capillary pressure.

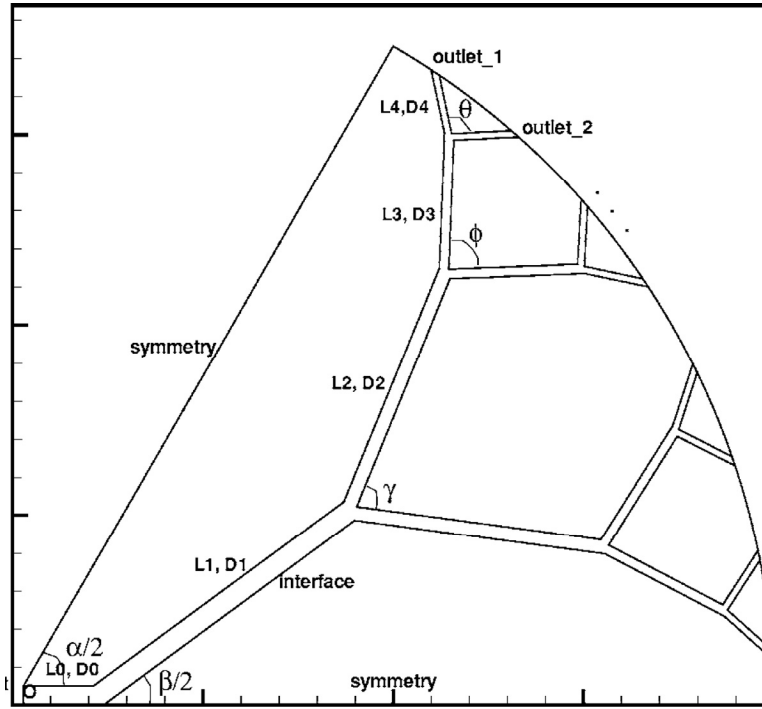


Fig. 1.11 – An example of a fractally branching channel layout designed to minimize pressure drop (from [1.13]).

1.4 Evaporator Optimization

As was hopefully demonstrated in the preceding sections, the μ C-LHP is a highly integrated, multifaceted system which involves many different and wide-ranging research areas, each no doubt worthy of its own thesis-level study: CPS fabrication, capillary flow modeling, transport channel design, nonstandard wafer bonding, charging and sealing, just to name a few. This particular study is focused almost exclusively upon maximizing the performance of the primary evaporator surface (*i.e.* maximizing heat transfer coefficients and overall vaporization rates from the “ceiling” of the Columnated Vapor Chamber). While this goal has remained more or less unchanged since the project’s inception, our team (and specifically the present author) has attempted to meet this challenge using two somewhat divergent approaches.

At the beginning of the project, my primary focus for optimizing evaporator performance was *nucleate boiling*. Nucleate boiling (see Section 2.4.1) usually provides the highest conventional heat transfer coefficients and has long been the industry standard for many macroscale boiler applications. Not surprisingly, there has been a sustained push to develop similar phase change systems for microfluidic heat exchangers [1.17]. Indeed, for many researchers, the “holy grail” of IC cooling would involve fully integrated microchannels *within* the device itself, through which liquid coolant would flow and boil in the nucleate regime.

Unfortunately, there are a number of phase-change phenomena unique to the microscale that make such systems as yet untenable. For a more detailed description of such phenomena, as well as a summary of how leading researchers have attempted to overcome the challenges posed by them, the reader is encouraged to see Appendix A, Sections III.2.2-4. For now, it suffices to say that, in general, nucleate boiling in microchannels tends to be terribly unsteady and prone to spontaneous failure. As Carey *et al.* [1.18, 1.19] have argued quite convincingly, this is due primarily to the tremendous thermal energy required to reactivate a “dormant” nucleation site (*i.e.* to reinitiate the ebullition, or bubbling cycle). Consequently, my first approach was to explore the feasibility of a novel surface treatment that would utilize passive *self-activating* nucleation sites to help stabilize nucleate boiling in microchannels. A brief summary of that research can be found in Chapter 3, while a complete version is included as Appendix B.

Although that study produced some interesting and promising results, for a number of reasons (enumerated specifically in Section 3.6), I shifted the focus away from nucleate boiling and back toward more traditional thin-film evaporation as the preferred vaporization mode for the μ C-LHP. As a result, my second approach (discussed in Section 1.3.3 above) was focused instead upon dramatically increasing the thin-film evaporation rate by creating a column-ceiling interface designed specifically to increase the area of the high-transport interline region. Although the physics involved in interline evaporation is still neither entirely understood nor precisely modeled, the final design of the μ C-LHP vapor chamber will be based primarily upon the *experimental* results obtained from testing a wide variety of column and evaporator layouts using the open-loop planar microevaporator that constitutes the bulk of this study (Chapters 5-8). Specifically, it is my belief that by properly sizing and spacing the columns and controlling the wetting characteristics of the interfaced evaporator surface, we can improve the overall evaporator performance by maximizing the amount of evaporator area whose nominal film thickness will lie in the interline region of the meniscus formed between adjacent columns. Even if the interline region cannot be maximized in a strict mathematical sense, simply increasing the area by a modest percentage will greatly increase the overall heat and mass transfer because the effects of the interline region are so dominant on total evaporation.

Chapter 2: Background on Phase Change Heat Transfer

This chapter is divided into four sections, all of which attempt to familiarize the reader with the basic heat transfer concepts essential to understanding the bulk of this study. In the first section, I introduce the concepts of thermal driving force and thermal resistance, as well as the classical heat transfer modes of conduction, convection, and radiation. In the second section, I discuss the phenomenon of phase change and its associated latent heat, particularly as it relates to overall heat transfer. Finally, in the third and fourth sections, I explain in much greater detail the liquid-to-vapor phase transformation (*i.e.* nucleation and boiling), which naturally forms the basis for many aspects of this research.

2.1 Elementary Heat Transfer Concepts

The primary focus of this study is heat and mass transfer, particularly the complex phenomena associated with liquid-to-vapor phase change. While the motivation, fabrication, and applications are clearly geared toward the MEMS community, the real “science” remains firmly rooted in energy transport. Consequently, Chapter 2 is designed to familiarize the “MEMS-centric” reader with the essential concepts of heat transfer that form the core of this study. As discussed in the Preface, some of the more general material will be relevant throughout the entire study, while some very esoteric topics will only be necessary for a thorough understanding of the Self-Nucleating Evaporator project summarized in Chapter 3. However, rather than simply repeatedly referring the reader to my Master’s Thesis, I have chosen to include within the body of this dissertation the necessary material to create a stand-alone document. It is my opinion that any study on phase change should provide a basic introduction to the jargon and governing equations of heat transfer, even if such information is readily available in other outside sources. In any case, much of the material from Sections 2.3 and 2.4 below is drawn from Chapter 5-7 of *Liquid-Vapor Phase-Change Phenomena* by Carey [2.1], as well as Professor Carey’s lectures at UC Berkeley (Mechanical Engineering 258: Heat Transfer with Phase Change).

In the most general sense, heat transfer involves the transport of thermal energy from one location to another, due to a temperature difference. Wherever there exists a finite temperature difference within a medium or between two media, thermodynamics demands that heat transfer must occur. However, while thermodynamics predicts only the *final energy states* of the system, heat transfer supplies the *rate equations* at which energy is transferred. In most engineering applications, the goal is typically to control the temperature of various locations in a system by controlling the heat transfer between them. Some applications attempt to maximize heat transfer; others attempt to minimize it. However, universal to all varieties and applications of heat transfer are the concepts of *driving force* and *resistance*. In much the same way that voltage is the driving force for the flow of electric current against various electrical resistances, temperature difference is the driving force for the flow of thermal energy against various thermal resistances. The three basic types of heat transfer are conduction, convection, and radiation. Conduction transports energy within stationary media (either solid or fluid), convection transports energy between a solid surface and an adjacent moving fluid, and radiation transports energy between two objects of different temperature via electromagnetic waves (which do not require intervening media).

2.1.1 – Conduction

The starting point adopted here for conduction heat transfer is Fourier’s law. Based on phenomenological evidence of one-dimensional conduction through a rod, it states that the total heat transfer (thermal energy passing through the rod) is directly proportional to the temperature difference across the rod, directly proportional to the cross-sectional area of the rod, and inversely proportional to the length of the rod. In other words:

$$q \propto -A \frac{\Delta T}{\Delta x} \quad (2.1)$$

Note the minus sign, which indicates that heat is always transferred in the direction of decreasing temperature (*i.e.* in the direction opposite the gradient). To transform this proportion into an equality, it is natural to define a thermal conductivity k as the constant of proportionality and to assume it to be an intrinsic property of the medium. If we now instead concentrate on the heat *flux* (thermal energy per unit area) and evaluate Eq. 2.1 in the limit as $\Delta x \rightarrow 0$, we obtain the more ubiquitous form of Fourier’s law:

$$q''_{cond} = -k \nabla T \quad (2.2)$$

This equation holds regardless of geometry and coordinate system, so long as the del operator is performed correctly. This equation easily communicates the basics of thermal energy transport, particularly when viewed from the “driving force vs. resistance” standpoint. The driving force for conduction is the temperature gradient within the medium, while the resistance to conduction is the reciprocal of the thermal conductivity k .

2.1.2 – Convection

The basics of convection are much the same. However now we have two media: the solid surface at one temperature and the adjacent moving fluid at another. While the details of convection are much more complicated, we can still construct a heat flux equation with much the same form:

$$q''_{conv} = h(T_{surface} - T_{fluid}) = h \Delta T \quad (2.3)$$

Here, the heat flux from the surface to the fluid (or vice versa) is once again directly proportional to the temperature difference between them, with the constant of proportionality h defined as the convection heat transfer coefficient. Though obtaining the value of h is usually a difficult task, we can once again communicate the basics of thermal energy transport. The driving force for convection is the temperature difference between the fluid and the surface, while the resistance to convection is the reciprocal of the heat transfer coefficient h .

2.1.3 – Radiation

Heat transfer by radiation is notably different from conduction and convection in that it does not require a medium to transport energy. As mentioned earlier, the thermal energy associated with radiation is instead transported by electromagnetic waves (or alternatively, photons) and occurs most efficiently in a vacuum. Radiation is an extremely complicated

phenomenon, but for this study it will suffice to explain the barest essentials. Basically, a surface (or gas, fluid, *etc.*) at finite temperature T_s will *emit* thermal energy E as

$$E = \epsilon \sigma T_s^4 \quad (2.4)$$

where ϵ is a radiative property of the surface known as *emissivity* (a measure of how effectively a surface emits energy; $0 \leq \epsilon \leq 1$), and σ is the Stefan-Boltzmann constant ($\sigma = 5.67 \times 10^{-8} \text{ W/m}^2 \cdot \text{K}^4$). At the same time, a surface will also *absorb* thermal energy G_{abs} as

$$G_{abs} = \alpha G \quad (2.5)$$

where α is another radiative property of the surface known as *absorptivity* (a measure of how effectively a surface absorbs energy; $0 \leq \alpha \leq 1$) and G is the sum of all incident radiation known as the *irradiation*. For the special (but very common) case in which a small *gray* object (“gray” in this technical sense means that $\alpha = \epsilon$) is surrounded by a much larger isothermal enclosure at finite temperature T_{sur} , the *net* rate of radiation heat flux *from* the surface can be approximated by

$$q_{rad} = E - G_{abs} = \epsilon \sigma (T_s^4 - T_{sur}^4) \quad (2.6)$$

Note that the driving force for radiative transport is not simply a “linear” temperature difference (as it was for conduction and convection) but is instead dependent on temperature raised to the fourth power. For this reason, radiation heat transfer can often be neglected in situations where the temperature difference is relatively small. For example, radiation will not play a significant role in this study at all.

In most heat transfer applications, several modes of heat transfer may occur at different locations as the energy is transported from its source to its sink. As a familiar example, consider the planet earth (and ignore the sun for now). Radioactive decay deep in the core can be considered the original heat source. This heat is first conducted through the solid core, then both conducted and convected through the viscous mantle, then conducted through the crust, then conducted and convected through the oceans and atmosphere, and then ultimately radiated away into deep space. All of these heat transfer modes have different driving forces and different resistances. Some are in series, others in parallel. How can we determine the overall heat transfer? Luckily, the electrical analogy discussed earlier becomes quite useful in this case. Just as we can construct networks of batteries and resistors to analyze the flow of electricity in complicated electrical circuits, so we can also construct networks of temperature differences and thermal resistances to analyze the heat flow in complicated multi-mode heat transfer systems (see Fig. 2.1). This will allow us to construct overall heat transfer coefficients and overall thermal resistances.




Fig. 2.1 – Construction of an electrical circuit analog to solve a multimode heat transfer problem.

2.2 Phase Change

Convection heat transfer may also involve phase change (*i.e.* vaporization and/or condensation of the adjacent fluid). Because phase change heat transfer is the primary focus of this study, it will be discussed in great detail in subsequent sections. However, for now it is worthwhile to note its influence on the “driving force vs. resistance” relationship. First, let us differentiate between *sensible* and *latent* heat. Sensible heat refers to thermal energy transfer which results in a change of *temperature*, while latent heat refers to thermal energy transfer which results in a change of *phase*. For example, suppose you drop an ice cube at -10°C into the ocean (at, say, 20°C). As the ice cube absorbs energy from the ocean, its temperature increases until the ice cube reaches 0°C . This is sensible heat (*i.e.* we can “sense” the heat absorption by measuring the temperature change of the ice). As the ice cube continues to absorb energy from the ocean, however, its temperature does not increase above 0°C . Instead, it begins to melt, breaking down the solid crystalline structure and undergoing phase change from solid to liquid. This is latent heat (*i.e.* the heat absorption is hidden, or “latent”, from our senses).

Because phase change heat transfer does not require an additional temperature difference, the absorption or release of latent heat allows for extremely efficient heat transfer in terms of the “driving force vs. resistance” relationship. In the case of convection, if the temperature difference between the solid surface and the fluid is sufficient to initiate a phase change, energy can be transported across the solid-fluid interface with very little resistance. In terms of Eq. 2.3, if ΔT is large enough to overcome the thermodynamic barrier of phase transformation (to be discussed later), the heat transfer coefficient h will become very large and greatly enhance the thermal energy transfer.

An additional benefit of phase change heat transfer is the isothermal nature of phase transformation. Recall the ice cube melting in the ocean. Because only latent heat is being released, the entire surface of the ice cube is held at a constant temperature (namely 0°C) while the ice cube changes phase. Analogously, if heat is transferred from a solid surface to a liquid by boiling the liquid, the temperature of the solid surface can be kept essentially uniform over a large area. This phenomenon is extremely valuable from an engineering standpoint, particularly if the characteristics or “performance” of a surface is dependent on temperature. For example, if the surface to be cooled consists of a large array of IC-based sensors (whose precise electrical properties, and thus resolution, depend on thermal homogeneity), then such isothermal cooling is absolutely essential. That, in essence, is the entire rationale for the DARPA TGP program. In the same way that an electrical ground plane creates a universal, board-wide “isopotential ground”, the TGP is an attempt to create a universal, board-wide “isothermal ground”.

2.3 Nucleation

Water does not simply start boiling the instant it is heated to its saturation temperature (*e.g.* 100°C at atmospheric pressure). *Saturation* merely forms the lower thermodynamic barrier below which phase change is not possible. At the other end of the spectrum, water cannot be heated past its *spinodal* temperature (see Appendix A, II.2.2) without violating mechanical laws of phase stability. Between these two temperatures, however, the fluid enters a metastable, *superheated* state. Superheat is defined as the amount the temperature is *above* the saturation temperature of the fluid (at the prescribed pressure). For example, at atmospheric pressure, water at 120°C would be said to have 20°C of superheat. In this metastable state, the fluid

density is poorly defined and its local “phase” is in constant spatial and temporal flux. During such *heterophase* fluctuations, microscopic vapor embryos spontaneously appear and disappear within a probabilistic framework. At one instant, more liquid molecules are evaporating than condensing, and the embryo grows; at another instant, more liquid molecules are condensing than evaporating, and the embryo collapses. Whether the embryo grows to the point at which liquid will “permanently” change to vapor (*i.e.* the onset of nucleation) is dependent on a number of factors including the liquid superheat, the size and distribution of the spontaneously generated vapor embryos, the surface tension of the liquid, and the thickness of the thermal boundary layer (if nucleation is occurring along a solid-liquid interface). I will not here derive the quantitative relations governing embryo growth and collapse (for that, please see Appendix A, II.3) but will instead simply state the primary consequences:

1. Excess energy is required to generate a phase interface. To generate a liquid-vapor interface large enough to support sufficient evaporation for sustained embryo growth, substantial superheat in the liquid is required.
2. Because of an exponential dependence on temperature, the theory predicts a very narrow range of temperatures for which spontaneous nucleation is favorable. Below this temperature, phase change is extremely unlikely; once this temperature is reached, phase change occurs essentially instantaneously.
3. For water at atmospheric pressure, this *kinetic limit of superheat* is approximately 300°C.

In other words, below 300°C, it is extremely unfavorable thermodynamically for liquid water to change to steam, whereas above 300°C, spontaneous change is practically assured. This naturally raises the question: Do I really need to heat a pot of water on the stove to 300°C to get it to boil? Not surprisingly to most readers, the answer is no. But the real question is, why?

Although the conclusions of the previous section are certainly legitimate, such high levels of superheat are rarely observed experimentally. For a pot of water on the stove, 10-15°C of superheat are typically sufficient to initiate nucleate boiling (once the bulk liquid has been heated to saturation). The source of disagreement can be found in our original idealization of a vapor embryo in contact with a solid surface. The presence of *insoluble gases* (not water vapor) at the surface was not considered. In reality, most surfaces are covered with naturally-occurring or machined scratches, pits, or other irregularities. Depending on their size and geometry (as well as the wettability of the solid), such cavities can trap insoluble gases when the surface is initially covered with liquid (see Fig 2.2). Because the advancing liquid front is convex (to maintain a consistent contact angle), the “nose” of the front will contact the opposite side of the cavity before the contact line reaches the bottom of the cavity (as long as $\theta > 2\beta$). How does this change anything? Remember, the extremely high levels of superheat discussed in the previous paragraph are only necessary to spontaneously *generate* a large liquid-vapor interface at which further vaporization occurs. With a trapped gas bubble, such an interface already exists! This naturally occurring liquid-gas interface allows vaporization to proceed at much lower superheat than would be required for “true” spontaneous nucleation of a new vapor embryo. Once the vaporization process is initiated, newly formed *vapor* can refill the cavity, allowing sustained low-superheat nucleation at the now-stable liquid-*vapor* interface.




Fig. 2.2 – a) The convex nose of an advancing liquid front will contact the opposite sidewall of the cavity if $\theta > 2\beta$. **b)** This will entrap insoluble gas to provide an initial embryo (after [2.1]).

As bubbles grow and detach from the surface (a process known as ebullition), a downwash of colder liquid from above breaks up the thermal boundary layer (see Fig. 2.3a). This lowers the liquid temperature near the surface and prevents bubble growth beyond the cavity. Transient conduction from the solid into the liquid eventually thickens the thermal boundary layer, providing superheat for further vaporization from both the solid surface and the surrounding liquid (Fig. 2.3b). Eventually, the bubble grows large enough such that its buoyancy and drag overcome surface tension and release it from the surface. During release, a small amount of vapor (and/or insoluble gas) is left within the cavity, providing the interface to allow the ebullition process to continue stably at low-superheat (Fig. 2.3c).




Fig. 2.3 – The ebullition process (after [2.1]).

Although the ebullition process is nominally quite stable, there are several ways in which nucleation sites can become deactivated (*i.e.* fully flooded). The simplest way occurs if the liquid entirely wets the surface to begin with. Referring back to Figure 2.2, if $\theta < 2\beta$ (*i.e.* for highly wetting fluids, extremely “liquid-philic” surfaces, and/or “obtuse” cavities), insoluble gases may not become entrapped within cavities at all. Highly pressurized systems may also

result in deactivated nucleation sites. Increasing the pressure increases the solubility of the ambient gas within the liquid and also drives liquid into cracks and cavities. Another common way that sites become deactivated is by exposure to subcooled liquid, causing complete condensation (and thus collapse) of the leftover vapor embryo. This can occur during cyclic boiling processes (when the heating cycle is “off”) or simply when the bulk liquid is purposely subcooled to increase heat transfer during boiling.

2.4 Boiling Curves

The next concept to consider is the actual heat transfer during phase change. A good starting point for a more graphical analysis is what is known as a *boiling curve*. It is a simple two-axis graph that quickly illustrates the relationship between the surface temperature at the solid-liquid convection interface and the resultant heat transfer. Surface temperature is traced along the x -axis, while the heat flux is traced along the y -axis. However, in order to “standardize” the curves for a wide range of fluids (which naturally require wide ranges of surface temperatures to produce various phase change phenomena), the curves do not plot *actual* surface temperature. Instead, they plot *excess* temperature or *superheat*. The higher the superheat, the more potential for heat transfer. Or, in terms of our “driving force vs. resistance” view, superheat is just the *normalized* driving force for phase change convection heat transfer.

Boiling curves can quickly communicate a tremendous amount of information based on the “driving force vs. resistance” view. For example, a steep, positively-sloped section indicates that a small increase in surface temperature can produce a large increase in heat transfer. Conversely, a near-horizontal section indicates that surface temperature can increase dramatically without a noticeable increase in the heat transfer. Thus the slope of the boiling curve provides a quick indication of the *resistance* to heat transfer. Furthermore, rapid changes in slope or discontinuities indicate *phenomenological* transitions, often from one boiling regime to another. Such regimes and transitions will be discussed more thoroughly in the next section.

One final issue that must be discussed is the actual method of constructing the boiling curves from experimental data. Determination of the independent and dependent variables is consequential. In particular, do we wish to measure the heat flux while we vary the superheat, or do we wish to instead measure the superheat while we vary the heat flux? Obviously, the two are inherently dependent, so it is impossible to control and/or measure both. Historically, boiling curves were constructed by electrically heating a platinum or nichrome wire passing through a large, quiescent liquid pool. The heat flux was *directly* varied by controlling the power supplied to the wire (*i.e.* by varying the voltage or current), while the surface temperature of the wire was *indirectly* inferred from its temperature-dependent electrical resistance. From a practical standpoint, such a “flux-controlled” boundary condition is much simpler to implement because it does not require any feedback circuitry to control the temperature. However, a “superheat-controlled” boundary condition is also possible, and perhaps surprisingly (as will be discussed later), the resulting boiling curve is actually quite different.

Boiling on the solid surface of an object immersed in an extensive, quiescent liquid is typically referred to as *pool boiling*, and it was such boiling conditions that were first studied using boiling curves. Though some additional criteria concerning heated area vs. bubble size, surface wetting, and surface roughness will be imposed later, at this point it is appropriate to introduce the various pool boiling regimes, as well as the overall boiling process.

2.4.1 – Superheat-Controlled Pool Boiling

We will start with a surface whose temperature corresponds to the saturation temperature of the fluid (*i.e.* superheat = 0). With no superheat, there is no driving force for heat transfer and consequently no flux across the solid surface (point *a* in Fig. 2.4). Now we begin to increase the temperature of the surface above the saturation temperature of the liquid. (Note that we are implementing a superheat-controlled boundary condition.) As mentioned earlier, because there is a certain thermodynamic barrier to phase transition (to be discussed in subsequent sections), low levels of superheat will not initiate a phase transformation. Instead, we enter the regime of single-phase convection heat transfer (path *a-b*). Heat is transferred from the slightly superheated surface to the liquid via the natural convection. Note that in this region the slope of the boiling curve is very slight. This indicates that single phase convection has a high resistance to heat transfer, or analogously, a low value of the heat transfer coefficient *h*.




Fig. 2.4 – Superheat-controlled boiling curve during heating (after [2.1]).

As we increase the surface temperature, the superheat will eventually become large enough to overcome the thermodynamic barrier for phase transition. At this point (point *b*), vapor embryos will nucleate at cavities along the surface, and these bubbles will then grow and detach from the surface. This is known as the onset of nucleate boiling (ONB). This additional *mode* of heat transfer increases the heat transfer coefficient *h* and thus leads to increased heat flux *without an increase in temperature*. Graphically, the boiling curve traces vertically at this point. This distinct slope transition indicates that we have entered the *nucleate boiling regime* (path *c-d*).

At first, the nucleation sites are few and widespread, and the boiling is said to fall in the *isolated bubble regime*. As we further increase the superheat, more and more nucleation sites are able to overcome the phase transformation barrier to become active, and consequently the bubble density gradually increases. Eventually the active sites become so dense that they merge prior to bubble detachment, forming columns of large vapor slugs that rise through the liquid. This is known as the *regime of slugs and columns*, and it is here that we find the highest heat transfer coefficients.

As the wall superheat is further increased, vaporization may occur so vigorously that liquid is unable to completely rewet the surface upon the departure of each bubble. Vapor patches begin to appear periodically over the surface, and evaporation of the liquid film between the surface and these patches may completely dry out the surface in places. Because a vapor gap presents much higher thermal resistance than liquid, the heat transfer coefficient begins to decrease at this point. Consequently, the overall heat flux achieves a maximum just prior to any dryout. This is known as the *critical heat flux* (point *d*). Beyond this point, areas of transient dryout will *decrease* the overall heat flux, even as the wall superheat is increased. This is known as the *transition boiling regime* (path *d-e*). This regime is characterized by large, local fluctuations in heat flux as various regions are dried out and rewetted.

If the superheat is further increased, the entire surface will eventually become blanketed by a thin vapor layer (point *e*). This is known as the *film boiling regime*. From this point on (path *e-f*), the heat flux will increase monotonically with the applied superheat until the surface itself melts. However, because the heat must now be convected, conducted, and radiated across a vapor gap, the heat transfer coefficient is lower than for nucleate boiling (note the shallower slope). Or, in terms of the “driving force vs. resistance” model, the increased resistance of the vapor film now requires much larger increases in driving force (superheat) for equivalent increases in flux.

Once a surface has been heated to the film boiling regime, gradually decreasing the superheat will cause the system to progress through each of the previous regimes in reverse order. However, experimental evidence indicates that heating and cooling paths have slightly different curves, particularly in the transition boiling regime and near the onset of nucleate boiling. This can be explained as follows. During transition boiling, the resultant heat flux is determined predominantly by surface wetting conditions, a phenomenon most easily characterized by contact angle. As we progress from transition to film boiling (*i.e.* during heating), the spreading vapor patches result in a *retreating* contact line. On the other hand, as we regress from film to transition boiling (*i.e.* during cooling), the collapsing vapor film results in an *advancing* contact line. For surface/liquid combinations for which there is significant difference in contact angle between advancing and retreating contact lines, divergent wetting characteristics will produce significant “flux hysteresis” in the transition boiling regime (compare path *f-g-h* to path *d-e-f* on Fig. 2.5).

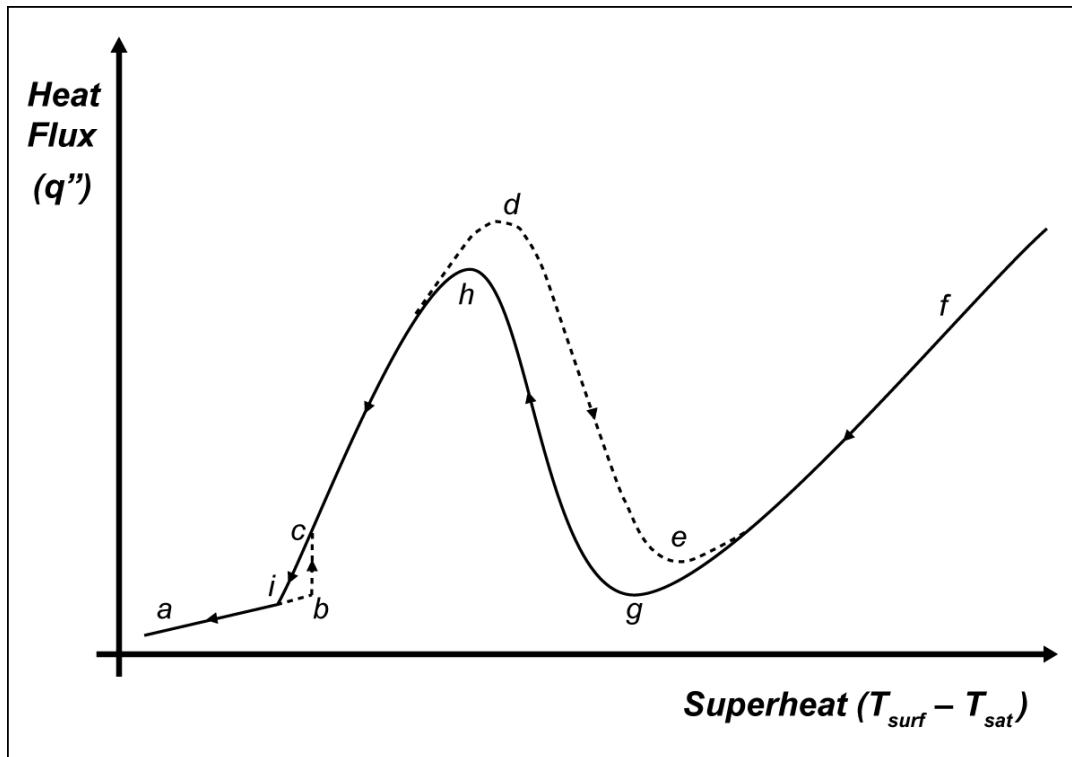


Fig. 2.5 – Superheat-controlled boiling curve during cooling (after [2.1]).

Additional asymmetry during cooling will often occur near the onset of nucleate boiling. During heating, nucleation sites are not activated until superheat is sufficient to overcome the thermodynamic barrier to phase change. However, once activated, these sites can remain active at superheat levels below those required for activation. Consequently, during cooling, rather than jumping vertically downward right at the ONB condition, the boiling curve will instead follow the nucleate boiling curve down along path *c-i* until it merges with the single phase natural convection curve along path *i-a*.

2.4.2 – Flux-Controlled Pool Boiling

We will once again start with a surface whose temperature corresponds to the saturation temperature of the fluid (*i.e.* superheat = 0, point *a* in Fig. 2.6), but this time, rather than controlling the superheat, we will instead gradually increase the heat flux. Although the single-phase convection regime of the boiling curve (path *a-b*) is unaffected by the change in boundary condition, something quite different occurs at the onset of nucleate boiling (point *b*). Recall that under superheat-controlled conditions, the additional mode of heat transfer provided by nucleation increased the heat transfer coefficient and resulted in a vertical jump in the heat flux. Under flux-controlled conditions, however, such a jump in heat transfer coefficient cannot lead to a corresponding jump in heat flux (because we are directly controlling it). Instead, there is a *decrease* in surface temperature to compensate for the *increase* in heat transfer coefficient, while the flux remains unchanged. Graphically, the boiling curve jumps horizontally to the left at the onset of nucleate boiling. In terms of the “driving force vs. resistance” model, increasing the




Fig. 2.6 – Flux-controlled boiling curve during heating (after [2.1]).

heat transfer coefficient decreases the overall thermal resistance, but since the energy flow itself is being regulated, we instead get a decrease in the required driving force (*i.e.* superheat).

Upon entering the nucleate boiling regime, the flux-controlled boiling curve once again coincides with the superheat-controlled curve (path *c-d*), and this congruence continues until the surface reaches the critical heat flux. Recall that at this point (point *d*), transient vapor patches on the surface begin to cause the heat transfer coefficient to *decrease*. In the superheat-controlled experiment, this naturally led to a decrease in the heat flux. In the flux-controlled experiment, however, we are *imposing* an increase in the heat flux. The only way to achieve this increase is to jump horizontally all the way to the film boiling curve (point *e*), skipping the transition boiling regime altogether. This sudden jump from nucleate to film boiling results in a tremendous increase in surface temperature. In many flux-controlled applications, the critical heat flux may be termed the *burnout heat flux* because of the potential to melt or otherwise damage components during this transition. Referring once again to the “driving force vs. resistance” model, the increase in resistance associated with vapor patches is countered by an increase in driving force (surface temperature, or superheat) in order to maintain the energy flow. Once the transition to film boiling is made, further increases in applied flux simply increase the wall superheat further, and the flux-controlled boiling curve again coincides with the superheat-controlled curve (path *e-f*).

From the film boiling regime, gradually decreasing the applied flux (starting along path f - g) will cause the system to progress through each of the previous regimes in reverse order. However, once the flux has been decreased to the minimum level required for stable film boiling (point g in Fig. 2.7), the curve must jump horizontally back to the nucleate boiling curve (point h) in order to further decrease the flux. Thus the transition boiling regime is bypassed altogether during cooling as well. Once into the nucleate boiling and single phase convection regimes (paths h - i and i - a , respectively), the flux-controlled cooling curve once again coincides with the superheat-controlled curve.




Fig. 2.7 – Flux-controlled boiling curve during cooling (after [2.1]).

Finally, let me again stress to the reader that this chapter has introduced only the barest essentials necessary to understand liquid-to-vapor phase change and heat transfer on a macroscale and qualitative level. For a deeper discussion regarding the thermodynamics of phase stability and the quantitative aspects of nucleation at the microscale level, please refer to Appendix A, Chapter II.

Chapter 3: Continued Development of a Self-Nucleating Evaporator Surface

This chapter attempts to summarize the research that constituted my first approach to optimizing evaporator performance in the μ C-LHP, namely stabilized nucleate boiling. Nearly all of this research is contained within my M.S. thesis, so the vast majority of the material has been relegated to Appendix B. However, additional experiments using a modified experimental setup were carried out *after* the completion of my M.S. degree, and I feel that the new results and conclusions are worthy of publication in this dissertation. I will therefore include within this chapter a comprehensive but highly abridged version of that study. It is my hope that this chapter alone will provide the reader with sufficient understanding of the basic concepts, testing methods, and results, but he or she is encouraged to consult the appendices for further detail and/or clarification.

The summary is divided into five sections. In the first, I discuss the challenges associated with microscale nucleate boiling and the motivation for a self-nucleating surface, while in the second, I introduce the concept of self-nucleation and develop a narrative explaining how such a concept could address some of those challenges. The third section describes the original experiments of my M.S. thesis and briefly summarizes the results. In the fourth section, I discuss the post-thesis experimental modifications and new results, while the fifth and final section briefly explains why I moved away from nucleate boiling for further μ C-LHP evaporator research.

3.1 Motivation for Self-Nucleating Evaporator

As briefly discussed at the conclusion of Section 1.4 above, there has been a great push in recent years to develop microfluidic heat exchangers that utilize phase change. A fantastic summary of motivations, past results, and a future “roadmap” for related research has been provided by Kandlikar [3.1]; here I will merely cite his primary conclusions:

1. The heat transfer coefficients in conventional flow boiling systems (> 3 mm hydraulic diameter) are very high compared to the corresponding single-phase values.
2. The required mass flow rates are reduced because of the use of latent heat in carrying the heat away, rather than just the sensible heat of the coolant (being limited by the available temperature rise in the coolant).
3. The heat removal process does not raise the temperature of the coolant as in the single-phase case, where the available temperature difference between the channel surface and the cooling fluid decreases along the flow length. This also results in a more isothermal evaporator surface, which may be critical for cooling highly temperature-sensitive electronics.

Unfortunately, the development of flow boiling microfluidic systems has been extremely limited by difficulties associated with 1) working fluid selection and 2) stable two-phase operation. With regard to the former, the low operating temperatures desired by electronics manufacturers preclude the use of water as a working fluid unless the system is operated at sub-atmospheric

pressures. At the same time, experimental (as well as ecological and health safety) data on other refrigerants is not yet sufficiently developed to effectively steer the designer toward the proper working fluid.

With regard to the latter difficulty, namely maintaining stable two-phase operation, there are many recent studies describing complicated flow phenomena that are unique to the miniscale and microscale. Perhaps the most commonly observed phenomena have been the extreme pressure fluctuations and backflow associated with rapid bubble growth. A 2005 experimental study by Kandlikar and Balasubramanian [3.2] provides excellent visual confirmation of these phenomena. Using high-speed photography (4,000-15,000 fps) of vaporization in parallel channels with a hydraulic diameter of 333 μm , they were able to observe nucleation along channel sidewalls, followed by extremely rapid bubble growth that quickly filled the entire channel width to form a large vapor slug. Interface velocities were as high as 3.5 m/s. Because of this rapid growth rate, the slugs do not “flow” along the channel (as in macroscale tubes), but instead fill the entire channel cross-section and expand along the channel in both directions (see Fig. 3.1 below). This naturally leads to backflow and enormous pressure drop fluctuations. Reverse slug growth even extended all the way into the inlet header in certain cases, causing severe flow maldistribution. A subsequent numerical simulation by Mukherjee and Kandlikar [3.3] confirmed these large interface velocities and backflow phenomena.




Fig. 3.1 – Violent vapor burst and backflow reported by Kandlikar and Balasubramanian (after [3.2]).

There have been a number of competing theories proposed to explain the phenomenological observations of microscale phase change that diverge dramatically from classical macroscale boiling models (see, for instance, [3.4, 3.5, 3.6]), but here I will simply refer the reader to Appendix A, Section III.2.3. One theory I wish to discuss in slightly more detail, however, is one put forward by Carey and others [3.7, 3.8] regarding spontaneous homogeneous nucleation at a deactivated site. As I found their model most convincing, it became the theoretical basis for my SNE concept.

In essence, they examined how a deactivated (fully wetted) nucleation site becomes reactivated. Recall from Section 2.3 above that it is the *generation* of a liquid-vapor interface that demands the large superheat values approaching the kinetic limit (e.g. 300°C for water at atmospheric pressure). If an interface is already present, further vaporization can continue at very low superheat. As a concrete example, consider a pool boiling scenario with an array of active nucleation sites exposed to a refrigerant with a saturation temperature of 34°C (see Fig. 3.2a). For stable nucleate boiling, a superheat of 10-15°C is typically expected, so we will assume a uniform wall temperature of 50°C. At this point, we then assume that a particular nucleation site becomes completely wetted and thus deactivated (Fig. 3.2b). Deactivation of the nucleation site eliminates one mode of heat transfer, which in turn leads to a localized hot spot beneath the deactivated site. This hot spot will grow and intensify until it develops sufficient superheat near the surface to initiate spontaneous nucleation. Local wall temperatures could easily reach 100°C, which would (at best) create large temperature gradients across the evaporator surface and could (at worst) damage or destroy underlying electronics. What's more, once a large enough interface has been nucleated (Fig. 3.2c), there is so much excess superheat in the surrounding liquid that vaporization at the liquid-vapor interface will proceed extremely quickly, resulting in violent vapor burst and possibly complete surface dryout (Fig. 3.2d). This can easily lead to the unstable backflow and failure conditions observed by Kandlikar. This very phenomenon became the starting point for my research. Namely, how can violent vapor burst during reactivation be avoided?

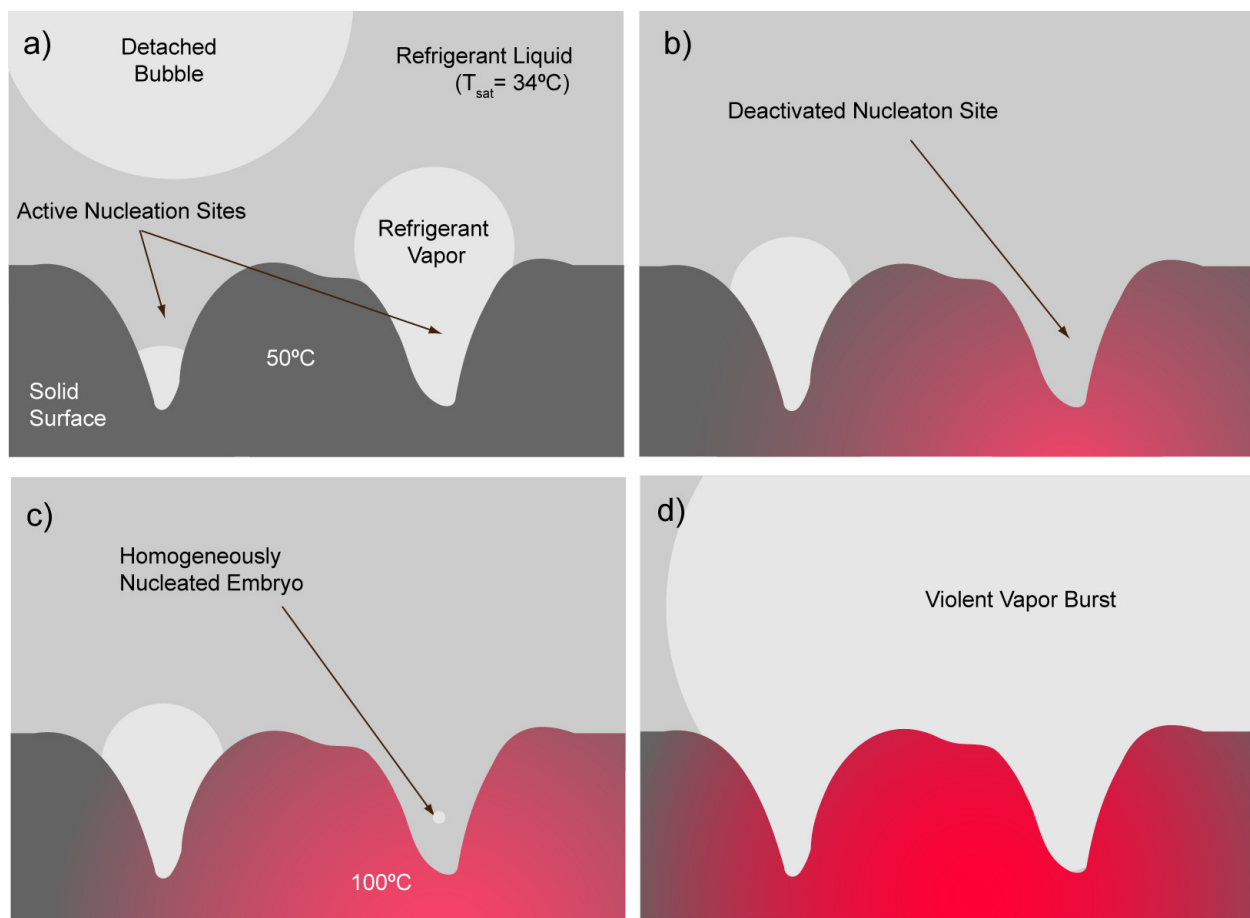


Fig. 3.2 – Sequence of site reactivation *without* self-nucleating surface treatment.

3.2 Self-Nucleating Evaporator (SNE) Concept

The guiding principle of this concept is the dramatic role played by entrapped gases during nucleation (see Section 2.3). Recall that insoluble gases entrapped within surface cavities allow nucleation to occur at much lower superheat because a gas/liquid interface is already present within the cavity (Fig. 3.3 top). However, if a cavity becomes completely wetted for any reason, site reactivation often requires large superheat (near the kinetic limit). This is why engineers have for years designed surface treatments that specifically inhibit complete wetting of nucleation cavities. Such techniques include the use of abrasive treatments, non-wetting coatings, porous coatings, and specifically designed reentrant cavity geometries [3.9 Ch. 11, 3.10, 3.11, 3.12 pp. 328-330]. For the SNE concept, rather than trying to prevent complete wetting in the first place, a *self-nucleating* surface instead *injects* insoluble gases into already-deactivated nucleation sites (Fig. 3.3 bottom), thus creating an “artificial” liquid-gas interface that allows vaporization to continue at low superheat.




Fig. 3.3 – Comparison between the natural entrapment of insoluble gas during wetting (top) and the injection of insoluble gas from a self-nucleating surface (bottom). Both create an artificial interface that precludes the high superheat required for spontaneous nucleation.

The key to any self-nucleating surface is the injection of insoluble gases into deactivated nucleation sites. While simple in theory, such a system requires not only an electro-chemical reaction of some kind, but also some method for determining which sites have become

deactivated. In addition, the system should be passive, self-regulating, and possess a functional lifetime appropriate for its application. For this study, I focused on the use of azobis-isobutyronitrile (AIBN), a chemical compound often used as an initiator for a variety of free radical polymerizations and also as a foamer in plastics and rubber. Its official IUPAC name is 2,2'-azobis(2-methylpropionitrile), with molecular formula $C_8H_{12}N_4$ and structure shown in Fig. 3.4. Its molar mass is 164.21 kg/kmol and its melting point is between 100-105°C. Physically, AIBN is a white, odorless crystalline solid, usually purchased as a fine white powder.




Fig. 3.4 – Molecular structure of azobis-isobutyronitrile (AIBN).

The most common chemical reaction associated with AIBN is its thermal decomposition. Like other azonitriles, the AIBN molecule is symmetric about the azo bridge (the double-bonded nitrogens), the most labile functional group. The azo bridge is easily thermally cleaved (at approximately 64°C), liberating gaseous nitrogen and leaving behind stabilized free radicals (see Fig. 3.5a). This reaction is the basis of the commercial utility of azonitriles as a source of free radical initiators and, to a lesser extent, as a source of nitrogen gas in foam blowing. According to the chemical equation, the AIBN will be reduced to 87% of its original weight once all of the nitrogen has been released. An experimental study by Hong *et al.* [3.13] showed that the temperature was approximately 100°C when the remaining weight reached 87% (see Fig. 3.5b), indicating that that nitrogen effluence occurred over the range of approximately 64-100°C. Upon further heating, a second thermal decomposition of the remaining free radicals occurs near the melting point.

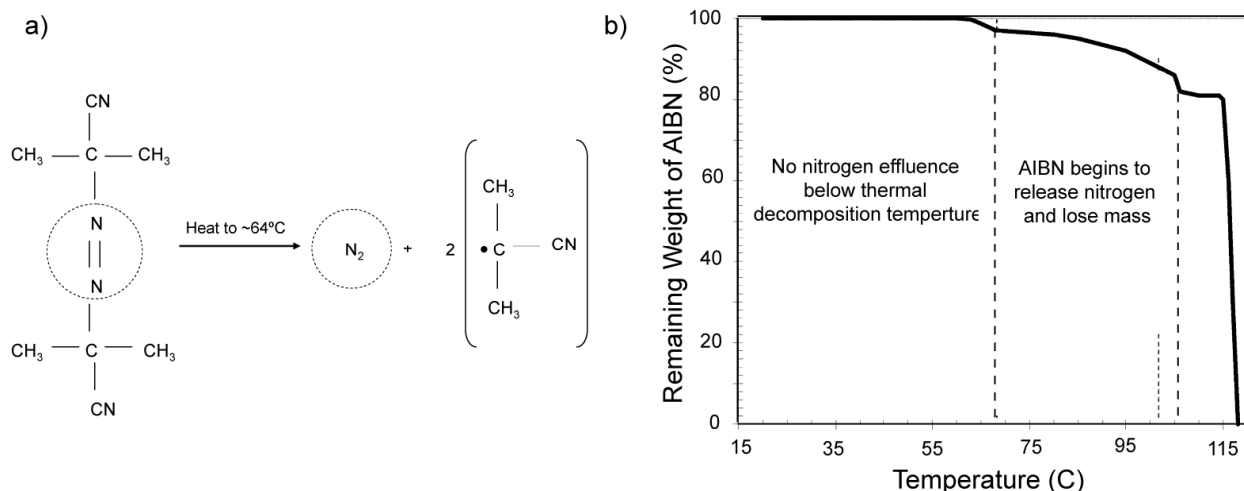


Fig. 3.5 – **a)** Thermal decomposition reaction of AIBN. **b)** Temperature range of decomposition (from [3.13]).

At this point, I would like to return to the site reactivation narrative discussed at the end of Section 3.1 and illustrated in Fig. 3.2. The only difference is that we will now assume that the surface is coated with a thin layer of encapsulated AIBN (see Fig. 3.6a). Just as before, deactivation of the nucleation site reduces the local heat transfer coefficient, which leads to an increase in wall temperature and the formation of an underlying hot spot (Fig. 3.6b). However, as this hot spot grows, it also heats up the AIBN surface directly beneath the deactivated site. Once the AIBN reaches its thermal decomposition temperature (approximately 64°C), it begins to release gaseous nitrogen, which then fills the cavity and creates an “artificial” entrapped bubble (Fig. 3.6c). This “injected” bubble provides a gas-liquid interface that allows further vaporization of the refrigerant to occur without the excessive superheat required to (homogeneously) nucleate a new embryo. As a result, we prevent the large temperature gradients and explosive vapor burst associated with the untreated site. Furthermore, controlled vaporization at the gas/vapor-liquid interface allows for a return to the normal ebullition cycle, which, in turn, raises the local heat transfer coefficient back to the original stable nucleate boiling level. With the original heat transfer coefficient restored, the localized hot spot disappears, the wall temperature decreases below 64°C , and the nitrogen effluence ceases (Fig. 3.6d). Thus the entire process is passive and self-regulating.

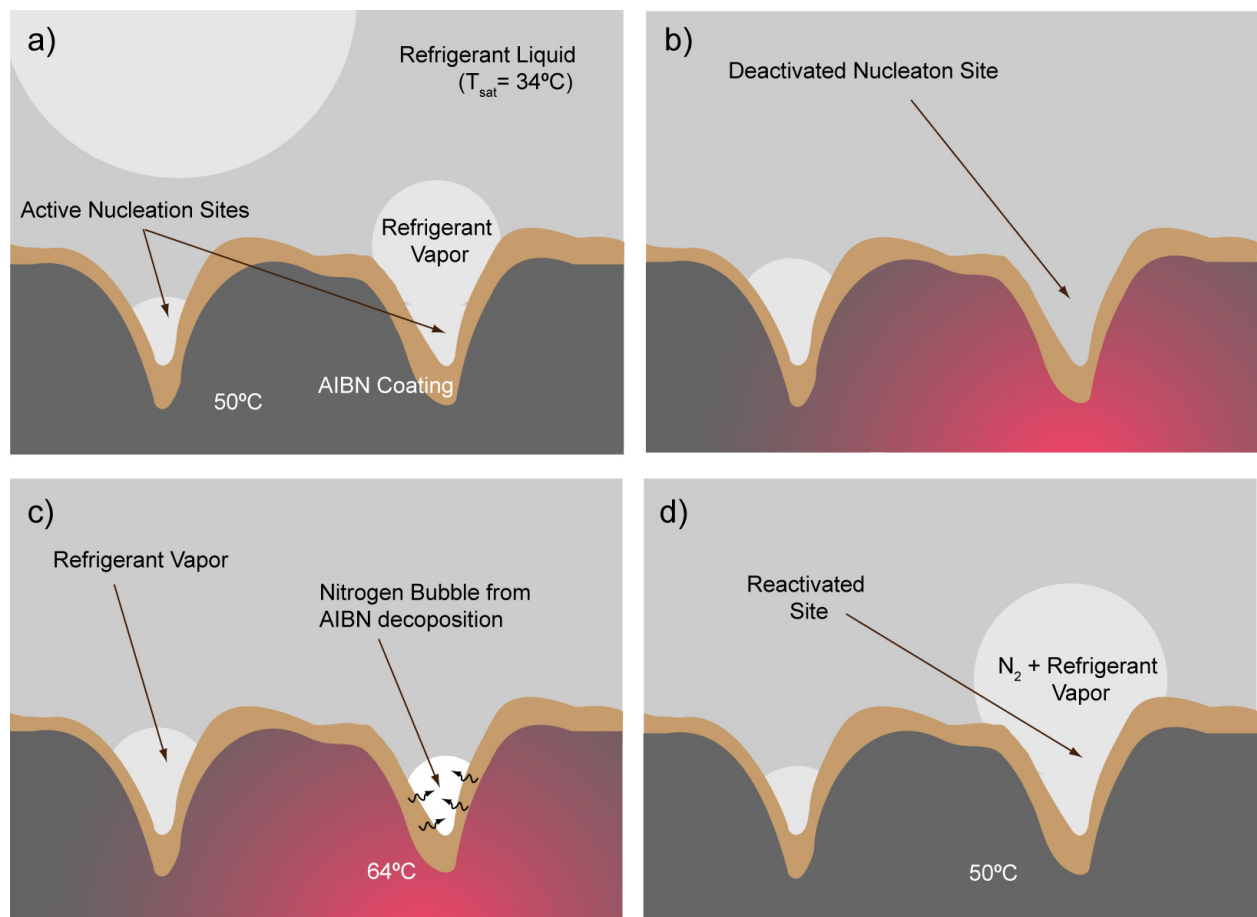


Fig. 3.6 – Sequence of site reactivation *with* an AIBN-based self-nucleating surface treatment.

3.3 Summary of Original Experiments

The goal of the experimental study was to determine whether an evaporator surface treated with an AIBN coating could significantly lower the characteristic superheat required to initiate nucleation during cyclic pool boiling. Although based on microscale phenomena, the experimental study itself was purely macroscale in nature, primarily to limit both cost and fabrication time (*i.e.* to eliminate dependence on the Microlab fabrication, for which I was not yet qualified). Moreover, the principal purpose of this study was to serve as a mere proof-of-concept for self-nucleation, providing an additional avenue for future enhancement of the TGP program.

The general experimental approach was quite simple. Take two surfaces with essentially identical physical characteristics (thermal conductivity, wetting properties, surface roughness, *etc.*), but beneath one surface, embed a thin layer of encapsulated AIBN. Next, expose both surfaces to a large pool of saturated refrigerant with a compatible boiling point, and then heat both surfaces from below with gradually increasing flux. As the heat flux is cyclically increased past the onset of nucleation and then decreased until nucleation ceases, continuously measure the temperature of both surfaces and look for variations in the maximum surface temperature between the AIBN and no-AIBN sides. In theory, deactivated sites with underlying AIBN will require much lower superheat to reinitiate nucleate boiling because nitrogen gas released into cavities will provide a gas-liquid interface for low-superheat vaporization without requiring homogeneous embryo formation. Such variations may not be seen during every heating cycle, of course, since only when a sufficient number of sites have become deactivated will there be a significant temperature rise over a measurable area of the surface. This is especially true during the initial cycles, when insoluble gases remain entrapped in cavities. At some point, however, the large superheat disparity between heterogeneous self-nucleation and traditional homogeneous nucleation should be reflected in the cyclic temperature maxima (see Fig. 3.7).




Fig. 3.7 – Expected qualitative results of cyclic boiling experimental study.

To test this concept, I constructed a submersible “diving bell” type test platform (see Fig. 3.8) that contained a variable-power ceramic heater embedded in an aluminum block. That block was spring-loaded against the “roof” of the diving bell and divided in half to create two identical heating surfaces. On the outer surface of the diving bell (directly above the two heated surfaces), the test surfaces were patterned, one with AIBN, one without. Thermistors were embedded in the heater block directly below the test surfaces, and these were connected to a data acquisition system to monitor surface temperature in real time. Finally, this entire test platform was submerged in the inner boiler of a double boiler system for testing (see Figs. 3.9 and 3.10). The outer boiler (a heated circulating water bath) was used to control the bulk temperature of the inner boiler, which was in turn filled with a variety of refrigerants. The ceramic heater within the aluminum block was used to heat the test surfaces from below, mimicking the dissipated heat from an integrated circuit in pool boiling conditions. (Note: for detailed information regarding fabrication and auxiliary systems, please see Appendix B, Chapter 5.)




Fig. 3.8 – Photograph and schematic pseudo-cross-section of submersible test platform.

Two different refrigerants were used in this study, HFE 7000 (scientific name methoxy-heptafluoropropane) and HFE 7100 (scientific name methoxy-nonafluorobutane). Both are Novec Engineered Fluids produced by the 3M company consisting of fully fluorinated hydrocarbon chains, with HFE 7100 possessing one additional fluorinated carbon atom. The first set of experiments involved HFE 7000 ($T_{sat} = 34^{\circ}\text{C}$), but unexpected boiling characteristics required a shift to a refrigerant with a significantly higher boiling point. Consequently, the second set of experiments involved HFE 7100 ($T_{sat} = 61^{\circ}\text{C}$). Both refrigerants are notable for their extremely low viscosity and surface tension and high dielectric strength (see Appendix A, Section V.4 for specific physical properties).




Fig. 3.9 – Schematic of double boiler experimental apparatus.




Fig. 3.10 – Photograph of entire SNE experimental setup.

A summary of the experiments is shown in Table 3.1. First, it should be noted that these experiments *did* display many characteristics supported by classical phase change heat transfer theory. For example, in all tests, the temperature required for the onset of nucleate boiling was higher than that required to sustain nucleate boiling once it had been initiated. Moreover, it was also observed that subcooling indeed raises the temperature required for the onset of nucleation and also raises the minimum temperature to which boiling can persist as the surface cools. As “highlighted” in my M.S. Thesis, however, the quantitative results of the cyclic pool boiling tests were extremely disappointing. *What I had hoped to observe in the comparative temperature-based data was not seen at all.* Namely, where I had expected to record dramatic differences in the temperature required for nucleation between the treated and untreated surfaces (see Fig. 3.7), I instead recorded essentially identical temperature curves throughout every boiling cycle.

Table 3.1 – Summary of boiling parameters during cyclic boiling tests

Refrigerant	Test Type	T _{bulk} (°C)	T _{bath} (°C)	Heating Time (s)	Cooling Time (s)	Number of Cycles	T _{ONB} (°C)	T _{DNB} (°C)
HFE 7000	Saturated Cycle	34	36	100	200	9	48	41
HFE 7000	Subcooled Cycle	28	29	150	100	5	75	70
HFE 7000	Endurance	30	30.5	40	260	160	50	44
HFE 7100	Saturated Cycle	61	62.5	15	585	20	64	61
HFE 7100	Subcooled Cycle	51	52	60	180	19	80	69
HFE 7100	Endurance	61	62.5	15	585	55	63	61

This was at first very perplexing. From visual observation, I could see that the surfaces treated with AIBN *did* reach the onset of nucleate boiling before the untreated surfaces, even if the thermistors recorded essentially identical temperatures. This led me to believe that it was not a problem with the self-nucleation *concept*, but rather with the heating and temperature measurement systems. Most significantly:

1. The large thermal mass and high thermal conductivity of the aluminum heating block caused tremendous difficulties with controlling heat flux. Prescribed, *symmetric* heat flux could not be delivered to the boiling surface quickly, and the block took an excessively long time to cool down between heating cycles.
2. Although their size, response time, and accuracy were more than sufficient for this study, the relative remoteness of the thermistors from the evaporator surface precluded even the crudest qualitative comparison-based tests. This was mainly due to the fact that they were embedded in the surface of the thermally-massive, highly-conductive aluminum heater block. Consequently, they consistently mirrored the steadily increasing and decreasing temperatures of the heater block, rather than the temperatures of the overlying boiling surface.

3.4 Modifications to Experimental Setup

Based upon the shortcomings of the experimental setup enumerated in the previous section, I made some relatively quick and simple design changes to try to improve my quantitative temperature data. First, although I was ultimately constrained by the relatively large dimensions of the macroscale ceramic heater, I trimmed away as much excess material from the aluminum heater block as possible. While clearly not nearly as ideal as microfabricated thin-film resistive heaters, it was a quick fix that reduced the thermal mass of the block by almost 50%. Secondly (and I believe more significantly), I bored out the aluminum directly beneath the thermistors and inserted nylon bushings to better thermally insulate the thermistors from the aluminum block (see Fig. 3.11). With this increased thermal isolation, it was my hope that the thermistors would more accurately measure the temperature of the overlying plastic surface, rather than the (largely uniform) temperature of the aluminum block.

In the initial experiments, the treated surfaces were coated with a slurry of AIBN in a primer-based polymer, while the untreated surfaces were simply coated with the polymer. For the second round of tests, to make absolutely sure that the AIBN-treated surface was not *structurally* superior to the untreated surface (*i.e.* possessing better nucleation cavities), I coated the untreated surfaces with a slurry of *pre-decomposed* AIBN (*i.e.* already heated to 100°C) in the polymer. In such a way, the two surfaces were structurally identical, but only the treated side would release nitrogen during heating. Finally, I was sure to thoroughly degas the refrigerant prior to testing by vigorously boiling the inner boiler for several hours.




Fig. 3.11 – Schematic of modified heater block with nylon inserts to better thermally isolate thermistors.

3.5 New Results

After making the aforementioned modifications to my experimental setup, I performed more or less identical cyclic pool boiling tests using saturated HFE 7100 ($T_{sat} = 61^\circ\text{C}$). The results are shown in Figs. 3.12 below. The red and pink curves indicate the temperatures recorded by two different thermistors embedded beneath the untreated surface, while the blue

and green curves indicate the temperatures recorded by two thermistors embedded beneath the AIBN-treated surface. Quite unlike my first experiments, the thermistors now indicated very different temperatures between the treated and untreated sides, confirming my conceptual prediction. For most of the cycles shown, the surface treated with AIBN achieved nucleate boiling with only 4-5 degrees of superheat, while the untreated surface reached temperatures 10-15 degrees higher (and did not even necessarily achieve nucleate boiling).

Note also that the effect of the AIBN began to wear off after about 11,000 seconds. This is in full concurrence with the finite nature of the AIBN reactants; eventually the thermal decomposition reaction must be exhausted. The unusually gradual nature of the AIBN decomposition is likely due to its encapsulation as a solid (as opposed to full dissolution in solution). While there is considerably less knowledge regarding the decomposition of AIBN in the absence of solvents, it is generally known to be an autocatalytic reaction below its melting point [3.14]. This not only means that one or more of the products are the same as one or more of the reactants, but often results in highly nonlinear (or even oscillatory) behavior.

However, what is more important for this particular application is the fact that the performance (as quantified by the surface temperature) degrades *gradually* and *observably*. Because the AIBN continues to release nitrogen from 64-100°C, the AIBN is not used up all at once, but is only progressively exhausted as the necessary nucleation reactivation temperature increases. In other words, once all of the nitrogen has been released for chemical equilibrium at 64°C, the AIBN *at that location* will simply heat to 65°C when it is again called upon for reactivation. Once chemical equilibrium at 65°C has been achieved, it will only need to heat to 66°C, and so on. This process can continue at a specific location (and, again, *only when deactivation necessitates it*) until the dissociation reaction ceases at approximately 100°C. Moreover, only a tiny amount of injected gas is required to reactivate (and then subsequently cool) a deactivated nucleation site. Normal operation and even occasional reactivation would not deplete the coating fast enough to render it impractical. This allows such a self-nucleating coating to serve as a performance monitor and extra level of safety against catastrophic failure by vapor burst and dryout.

Thus, in conclusion, the effluence of an insoluble gas from inactive cavities can immediately reinitiate nucleate boiling with little to no superheat, and the feasibility of a self-nucleating evaporator has been confirmed.




Fig. 3.12 – Top: Saturated HFE 7100 endurance test using modified experimental setup ($T_{\text{bulk}} = 61^{\circ}\text{C}$); **Bottom:** Zoomed-in view of a single cycle.

3.6 Rationale for Discontinuing SNE Research

Despite the success of my modified pool boiling experiments, there are a number of reasons why we have moved away from nucleate boiling toward thin-film evaporation for further research on the μ C-LHP evaporator. These include the following:

1. It is not clear whether capillary driven mass flow from the CPS wick will be sufficient to provide enough liquid for nucleate boiling.
2. Insoluble gases tend to greatly degrade the performance of closed-loop heat pipes. Dissolved nitrogen from reactivated sites would have to be gettered from solution in the condenser region. This would involve completely new structures and materials and greatly complicate our design.
3. Although (ideally) vapor would not be transported in-plane in “bubble” form (*i.e.* vapor completely surrounded by liquid), there is always the possibility of unexpected vaporization behavior resulting in channel blockage.
4. Microscale deposition of AIBN (screen-printing) still not developed or compatible with high-temperature anodic bonding and would add several levels of complexity and difficulty to the fabrication process.
5. Without an “externally” driven flow (*e.g.* dP/dx), thermally-imposed surface tension gradients would promote the migration of surface vapor bubbles *against* the desired direction of heat transfer (due to the Marangoni effect).
6. The fixed AIBN decomposition temperature would place additional constraints on the working fluid (for this study) or on the range of operating temperatures (for a closed-loop system). No doubt other compounds exist with widely ranging decomposition temperatures, but again, time does not permit this level of broad exploration.

This is not to say that self-nucleation is without merit or practical application. While the time and technology constraints facing the μ C-LHP project do not permit me further study on this topic, I see a wide variety of heat transfer applications for SNE, both at the macro- and microscale.

Chapter 4: μ C-LHP Hermeticity Studies

One peripheral but crucial aspect of the μ C-LHP project to which I dedicated significant time and effort was the development and testing of a variety of sealing techniques to satisfy the hermeticity milestones of the TGP program. This chapter summarizes the hermeticity experiments performed to fulfill DARPA's Phase I requirements, which ultimately contributed to our team's procurement of Phase II funding. The chapter is divided into four sections. In the first, I briefly discuss the motivation for hermetic packaging and the specific TGP program goals for hermeticity, focusing particularly on the overall μ C-LHP architecture and the associated sealing strategies and difficulties. In the second section, I describe the anodic bonding process (used to fuse the multiple layers of our device) and discuss its effect on hermeticity. Finally, the third and fourth sections discuss two different fill-port sealing techniques that were used (epoxy-based and solder-based), including the testing methods and results. In presenting this research, I first provide a summary of the sealing results, with particular emphasis on the final successful approach, and then describe the experimental process that led to these final conclusions.

4.1 Motivation for Hermetic Sealing

Leakage is essentially a "physical" phenomenon; it is a consequence of individual fluid molecules (liquid, vapor, or gas) working their way through the molecular structure of the encapsulating solid, driven by a concentration or pressure gradient. In the microelectronics industry, water vapor permeability especially has been studied extensively beginning in the 1970s, as moisture-based corrosion of electronic interconnect became a primary failure mode for many long-term encapsulated devices. The growth and maturity of MEMS technology (particularly resonator-based devices) generated further research on gas permeability, as device performance became highly dependent on effective vacuum packaging. Stroehle [4.1] was responsible for one of the most thorough early studies of microelectronic packaging permeability and acceptable leakage rates, and his results were used to generate one of the most easily-recognized figures on the subject (see Fig. 4.1).

While cracks and other structural defects certainly play a large role on permeability at a macroscopic level, at the microscale, the rate at which fluids diffuse through a solid is largely dependent on the density of the solid, the physical size of the liquid/gas molecules, and the pressure (or concentration) difference across the solid. More specifically, the denser the material (or the larger the fluid molecules), the lower the leakage rate. For this reason, truly hermetic systems tend to be metallic with welded seals, and standard industry leakage testing is performed using small-radii gases such as helium. For hermetically-sealed systems requiring electrically insulated feedthroughs, glass is often used alongside exotic alloys designed to match the thermal expansion coefficient of the glass (*e.g.* Kovar).

Heat pipes must maintain the proper amount of working fluid for a particular operating condition. Fluid leakage out of (or ambient gas leakage into) the closed-loop system will greatly degrade its performance. Consequently, existing macroscale heat pipes are usually fully encased in metal with welded seals and generally regarded as perfectly hermetic (*i.e.* zero leakage over the practical life of the device). Unfortunately, the thermal expansion matching requirements of the TGP program greatly limit the use of traditional metals as a structural material, so one major

challenge of the TGP program is demonstrating that non-metallic materials can also prevent fluid loss during extended operation. An additional complication for hermetic sealing is the highly planar geometry required. Because leakage is a physical phenomenon, increasing the thickness of the solid encapsulation (thus lengthening the diffusion path) will decrease the leakage rate, but here we are greatly limited due to overall thickness constraints. The specific hermeticity milestones for the various TGP phases are summarized in Table 4.1 below.




Fig. 4.1 – Water permeability through various solids (after [4.2]).

Table 4.1 – DARPA TGP Program Hermeticity Milestones

	State of Art	Phase I	Phase II	Phase III
Fluid Loss Per Year at 100°C	0%	10%	1%	0.1%

The primary structural material of the μ C-LHP is single-crystal silicon, which has very low permeability and thus excellent hermeticity (see Fig 4.1). Anodic bonding *between* silicon layers (see next section) generates a thin interfacial layer of silicon dioxide (glass), but glasses, too, provide very low permeability. Thus both our structural materials and our primary fabrication/package techniques are inherently compatible with stringent hermeticity requirements. The real obstacle for our design (particularly in the early phases) is sealing the fill-ports through which working fluid is introduced (see Sections 4.3 and 4.4), particularly if this is to be done at non-atmospheric conditions.

4.2 Anodic Bonding

Nearly all microfluidic devices require some sort of bonding to fully enclose the various fluidic channels and reservoirs. Temporary clamping or capping with a removable cover slide may be practical for research and testing, but any practical functional device is likely to be permanently capped using a nominally irreversible process. Microfluidic devices that utilize polymer substrates (PDMS or PMMA, for example) are often capped with glass wafers or other polymer substrates by some form of “gluing”. Usually this involves either partially dissolving a thin layer of one of the substrates prior to contact or briefly heating the clamped pair past their melting or glass transition temperature). For silicon substrates, however, other bonding methods are needed.

One of the strongest and most popular wafer-level bonding techniques for silicon is anodic bonding. In this process, a silicon wafer is permanently bonded to a glass wafer using high voltage and elevated temperatures. The glass wafers consist primarily of silicon dioxide but also contain trace elements (notably boron and sodium) to both increase electrical conductivity and match the coefficient of thermal expansion (CTE) of silicon. Such glasses are commonly called borosilicates, with Pyrex, Borofloat, and Dupont 7740 being perhaps the most well-known commercial brands. During the anodic bonding process, a thin layer of silicon is transformed into silicon dioxide right at the glass-silicon interface, essentially extending the glass wafer *into* the silicon. The physics behind this can be explained as follows (see Fig. 4.2 and [4.3]). With the silicon and glass wafers in good physical contact, a large electrical potential (several hundred to several thousand volts) is applied across the glass-silicon wafer stack, generating a strong electric field in the poorly-conductive glass. The highly mobile sodium ions (Na^+) in the glass are driven away from the interface toward the negative electrode (Fig. 4.2, left). This leaves the interface region depleted of Na^+ ions and oversaturated with less mobile oxygen ions. The excess oxygen ions then diffuse across the interface and react with the silicon to form an amorphous layer of silicon dioxide. This creates a permanent covalent bond with excellent strength. Generally, this process is performed at elevated temperatures (300-500°C) to increase both the ion mobility and the silicon oxidation rate, thus decreasing the necessary bonding time.

Note that two configurations are possible. When using *negative polarity* (Fig 4.2, top), the top electrode (cathode) contacts the glass at a single point, creating a non-uniform electric field. Since the field is strongest directly under the electrode, a point bond forms that then spreads outward over the surface of the wafer. While this tends to produce a higher quality bond (*i.e.* fewer voids), it generally requires a longer bond time. When using *positive polarity* (Fig. 4.2, bottom), the chuck (now the cathode) contacts the entire glass wafer, creating a more uniform electric field. Bonding occurs over the entire surface of the wafer simultaneously, often resulting in a poorer quality bond but requiring a shorter bond time.

During the bonding process (which typically lasts minutes to tens of minutes, depending on the process temperature and the wafer-to-wafer contact area), the current passing through the wafer stack usually follows a curve similar to Fig. 4.3a. This can be explained as follows. Initially (that is, upon application of the electric field), the current ramps up very quickly because the bond area grows, creating a larger cross section for electrons to pass. However, as more oxygen atoms diffuse across the interface into the silicon, the newly-generated silicon dioxide layer thickens. This thickening dielectric not only increases the electrical resistance across the wafer stack, but also creates a diffusion barrier between the oxygen-rich interface and the un-

oxidized silicon. Often anodic bonding tools have a current limiter to prevent shorting or arcing in the case of improperly loaded wafer stacks. In this case, the voltage quickly ramps up until the maximum current is reached, but remains short of the specified voltage. Then, as the silicon dioxide layer grows and the electrical resistance increases, the voltage can slowly increase to its maximum specified value while limiting the current to I_{\max} . This results in a transient flat-line saturated current curve (see Fig. 4.2b). Once the maximum specified voltage is reached and the silicon dioxide layer continues to grow, the current decays exponentially as in Fig. 4.2a.




Fig. 4.2 – During anodic bonding, a large electric field in the glass draws sodium ions away from the silicon-glass interface. This leaves behind excess oxygen ions, which then diffuse across the interface to form a thin layer of amorphous silicon dioxide. Note: the arrows in the right-hand images (after [4.3]) indicate the force felt by *electrons* (i.e. they depict $-E$, not E).




Fig. 4.3 – Characteristic current behavior during anodic bonding **a)** without current limiting; **b)** with current limiting.

Anodic bonding is rarely perfect, and small dust particles between the bonded layers may result in small voids in the glass interface (see Fig. 4.4). However, total system leakage can be largely mitigated simply by increasing the width of the bonded “border” region surrounding the functional device area (see top-down view in Fig. 4.5). A wider strip of anodically bonded silicon creates a longer in-plane diffusion path to inhibit leakage, and unlike device thickness (which *is* strictly limited in the TGP program), there is really no limit to the lateral extents of the device. In other words, the width of this barrier can be increased to the point where it is no longer the “weakest link” with regard to hermeticity.




Fig. 4.4 – Anodically bonded wafers, with a magnified view of a void. The “Newton’s Rings” are caused by the varying air gap thickness within the void.




Fig. 4.5 – Although anodic bonding may result in voids within the silicon dioxide interfacial region (dark gray), the diffusion path between the fluid-filled volume (light gray) and the outside environment can be sufficiently lengthened simply by increasing the bond width.

The specific tool used for anodic bonding in our research was a Karl Suss SB6 Thermocompression and Anodic Bonder (see Fig. 4.6), shortened to KSBonder in Berkeley Microlab nomenclature. The KSBonder is a semiautomatic, computer-controlled, stand-alone substrate bonder equipped with a vacuum/pressure chamber and a loading arm. It can process aligned or unaligned wafers (4" or 6"), as well as die-sized chips. The alignment accuracy is listed at 3 μm . Bonding pair alignment is done using the Karl Suss BA6 tool in conjunction with the Karl Suss MA6 Mask Aligner (KSAligner), although no precision alignment was needed for our experiments (since we were simply capping a patterned silicon wafer with an unpatterned Pyrex wafer). Vacuum capability is 5.0×10^{-5} Torr, but our bonding was done at atmospheric conditions. Lastly, maximum bonding pressure (primarily for thermocompression, not anodic bonding) and chuck temperature are 60 psi and 500°C, respectively. Most of our bonds were performed with zero bond pressure at between 300-450°C and 800V potential using negative polarity (glass on top of silicon). Current was limited to 10 mA.



Fig. 4.6 – Karl Suss SB6 Thermocompression and Anodic Bonder.

4.3 Port Sealing Summary

Having discussed the hermetic suitability of both our structural materials and bonding techniques, the next (and more challenging) issue is port sealing. For reasons discussed in this chapter's opening section, it was obvious to us from the beginning of the project that epoxy would not be suitable as a port-sealing material in any final version of a working device (*i.e.* a device that meets the TGP Phase III milestones). Generally speaking, they are simply not dense enough and are also quite unlikely to survive rigorous thermal cycling. On the other hand, the Phase I hermeticity milestone (<10% fluid loss over one year at 100°C) was fairly lenient, and we believed it could be satisfied using simple epoxy-based techniques already familiar to our lab. Moreover, we felt it was more prudent to dedicate the bulk of our time and energy to developing the more novel and fundamental aspects of our design (as well as meeting the other milestones), rather than trying to become experts on non-standard soldering techniques. After all, it wouldn't do us much good to meet the Phase III hermeticity milestones using some highly-advanced sealing method if we didn't meet the other Phase I milestones and thus failed to progress to Phase II funding. However, to demonstrate to DARPA that we were also looking ahead toward future phases with stricter hermeticity demands, we performed some brief studies with metallic solder-based ports seals as part of our Phase I demonstration, as well. Soldering microfluidic interconnect for hermetic sealing is quite nonstandard, and I must say that I found myself incredibly humbled by the difficulties it presented. Nevertheless, we were able to obtain some useful results.

Having just provided a general overview of the motivations guiding my port sealing research, I will now discuss the work in more detail. However, rather than discussing the entire process in chronological order, I will instead present the final results and techniques first, followed by a more detailed description of the trials and tribulations that led to these conclusions. First, as a guide to the reader, a summary of port sealing results:

- 1) Sealing the fill ports will almost certainly be the “weakest link” in the hermeticity game. It is the least automated, the least integrated, and the only “bonding/sealing” step that occurs in the presence of working fluid.
- 2) Epoxy-based sealing methods, the best results of which came using E6000 industrial sealant, were sufficient to meet the DARPA Phase I milestones, but cannot be used for later phases.
- 3) Because the “successful” E6000-sealed devices did not explode at elevated temperatures (unlike the solder-sealed devices), it is likely that expanding air (as well as some water) was able to diffuse through the seals during heating. However, the diffusion resistance (due mostly to the width of the seal) must have been sufficiently strong to prevent unacceptable fluid loss (at least for the given concentration difference across the seal).
- 4) Solder seals are very difficult to produce with regularity (at least manually by a novice), but they show excellent promise for effective later-phase DARPA hermeticity milestones.

4.3.1 – Solder-based Sealing Technique

The most promising sealing technique for future phases involves entirely covering a fill port with molten solder, which then adheres to the port mouth upon cooling. This seemingly simple task is complicated by a number of factors. First, solder does not adhere to Pyrex (glass) or silicon. If a drop of molten solder falls onto glass or silicon, it balls up (*i.e.* does not wet the surface) and simply falls off upon cooling. To improve adhesion, we had to first coat the Pyrex wafers with a thin layer of gold in the region where the solder would be applied (*e.g.* around the fill ports). However, like most microfabrication processes, even putting down a simple gold pattern requires several intermediate steps. The entire process, commonly called “lift-off”, is shown schematically in cross-section in Fig. 4.7 below. (Note: I will assume that the reader is somewhat familiar with microfabrication and will not attempt to describe all of these processes in detail. For the lay reader, a good starting point is [4.4, 4.5] .)




Fig. 4.7 – Cross-sectional view of metalized fill-port fabrication, including liftoff process.

The bonded Pyrex-silicon wafer stack is first cleaned in Piranha and then coated with several microns of photoresist. The resist is then patterned, exposing bare Pyrex regions where the gold will be put down. Unfortunately, gold by itself does not adhere well to Pyrex either, so first a thin layer of chromium (~50 nm) is sputtered onto the surface. Then approximately 150 nm of gold is sputtered on top of the chrome. These thin metallic coatings cover the entire wafer. When the wafer is put into acetone and agitated using ultrasonic vibration, the photoresist dissolves, taking with it the overlying chrome and gold (hence the name “lift-off”). This leaves

the chrome-gold layer covering only the regions where the pattern is desired (*i.e.* where the photoresist had been removed). At this point, the devices can be diced as before, and holes can be bored through the gold layer and underlying Pyrex into the reservoir to create fill ports. Figure 4.8 shows two complete devices with three metallized (but undrilled) fill ports, along with close-up views of two devices with the fill ports already drilled, forming the characteristic “bond ring”. Note that I tried two different metallization patterns (circle and 12-point star) to see if one might lead to better adhesion of the solder (*i.e.* meniscus-geometry-based surface wetting enhancement).




Fig. 4.8 – Top: Two undrilled devices with metallization over the fill ports. **Bottom:** Zoomed-in view of metallized ports that have subsequently been drilled for filling.

A second complication of solder-based sealing is the high temperatures associated with molten metal. Although solders are designed to melt at low temperature compared to other metals, the melting temperature is still far above the saturation temperature of water. This means that the molten solder ball will vaporize some of the water when it contacts the hole, effectively causing a small amount of steam to pour *out* of the hole at the same time one is trying to force the solder ball *into* the hole—a difficult task indeed. To make matters worse, the solder wants to stick *only* to the gold, so it was extremely challenging to keep the solder ball centered over the hole and completely covering it. Moreover, one typically had just one single chance to position the solder ball correctly because the gold tended to bond irreversibly with the solder ball upon

cooling. If one re-melted the solder to reposition it, it pulled off all of the gold coating and would no longer stick to the now-bare Pyrex.

Through much trial and error (and recognizing that two hands is not enough), I eventually developed a soldering technique that worked reasonably well (see Fig. 4.9). To help localize the heating and thus minimize water evaporation during soldering, I clamped the filled device to a large aluminum block, which served as a cold sink. I had one soldering iron mounted on an improvised z-axis controller (pulled from an old microscope) and positioned it over the edge of the gold bond ring, but not yet touching it. I then melted a large ball of solder on a second hand-held soldering iron and waited for all of the flux to burn off. (Note: I found this flux burn-off to be necessary because outgassing flux within the actual solder seal created large voids and prevented a hermetic seal.) Once this burn-off had occurred, I carefully lowered the z-axis controller such that the tip of the soldering iron contacted and heated the gold bond ring. The insulating Pyrex and underlying cold sink prevented most of this heat from reaching the water. After a few seconds, I let the large solder ball contact the gold bond ring and then raised the z-axis controller. Ideally the ball was large enough and placed such that it completely covered the hole while maintaining an unbroken contact line with the gold bond ring. Figure 4.10 shows a filled device with solder seals, along with close-up views of the sealed ports of two other devices.




Fig. 4.9 – Top: General setup and procedure for solder-based sealing. First the gold bond ring was preheated using the z-axis controlled soldering iron (1). Then, the molten solder ball was brought into contact with the gold bond ring (2). Finally, each iron was removed in reverse sequence (3 then 4), leaving a hardened solder ball completely covering the hole.




Fig. 4.10 – A solder-sealed device and two close-up views of solder-sealed fill ports.

4.3.2 – Significant Results

As discussed above, producing a successful solder seal was extremely difficult, and producing them with any regularity was next to impossible. However, I was able to seal a few devices for testing at elevated temperature. When these samples were placed into the 100°C test chamber, however, one of two things happened. Either they lost all of their fluid in the first day (indicating that the seal itself was bad) or else they simply exploded, shattering the Pyrex. While initially these explosions seemed like a terrible result, they actually proved that the seals on those devices were actually very, very good. Rather than leaking the expanding air (as the epoxy seals must have done), these seals did not allow any fluid (liquid or gas) to escape, thus creating a pressure build-up which eventually blew out the Pyrex. This was evident from the fact that the Pyrex always ruptured along the sidewall adjacent to the largest open membrane. Membrane deflection due to internal pressure is largest at the center of the membrane and subsequently generates the highest stress where the membrane attaches to the channel wall. For all but the Serpentine layout (see Fig. 4.18), the ruptures occurred along the edges of the large open cavities (see Fig. 4.11a, especially the topmost device). Likewise, the Serpentine device ruptured along the edges of its fill port reservoir, which constituted its own largest membrane (see Fig. 4.11b). To prevent the Pyrex from rupturing, I needed to effectively increase the stiffness of the membranes. Rather than change my wafer-level layout to decrease the size of the largest membrane (by inserting support posts or channel dividers, for example), I decided to instead just use a thicker Pyrex wafer. To first order, for an elastic membrane covering a pressurized cavity,

maximum deflection and stress are inversely proportional to the thickness of the membrane cubed and squared, respectively, that is:

$$\Delta z_{MAX} \propto \frac{1}{t^3} \quad (4.1)$$

$$\sigma_{MAX} \propto \frac{1}{t^2} \quad (4.2)$$

where t is the thickness of the membrane. For the next set of devices, I used a Pyrex wafer that was twice as thick as the original (1000 μm compared to 500 μm), and with this change, the devices survived the elevated temperature without rupturing. Even better, the ones that didn't leak immediately did not show any measurable fluid loss over the entire time I tested them, further supporting the fidelity of solder seals.




Fig. 4.11 – Solder-sealed devices that exploded during testing at elevated temperature. The small piece in the upper right of **(b)** is the blown out Pyrex reservoir membrane with the solder ball still attached.

4.4 Chronological Port Sealing Development

Having described the final solder-based sealing techniques and discussed the most important results, I will now detail chronologically the entire process that led to those conclusions, as well as provide a description of the device filling process and actual hermeticity testing procedures. As discussed earlier, although we recognized that epoxy-based sealing methods would not be suitable for a truly hermetically sealed final-phase device, we believed that epoxy-based methods would prove sufficient to satisfy the Phase I milestones; this indeed proved to be the case.

4.4.1 – Device Preparation and Filling

The layout of the Phase I device is shown schematically in Fig. 4.12a, while an undiced wafer-level photograph of four devices is shown in Fig. 4.12b. Note that the channels (and thus the fluid flow) in the Phase I device are strictly in-plane (as opposed to the out-of-plane columnar flow for the final μ C-LHP design described in Section 1.2). The single silicon structural layer has five basic components: 1) a small square evaporator/wick region consisting of very finely spaced parallel channels; 2) a large open square area that serves as the condenser; 3) vapor transport channels (to transport vapor from the evaporator to the condenser); 4) liquid transport channels (to transport condensed liquid from the condenser back to the evaporator); and 5) a liquid reservoir to interface with a filling port.

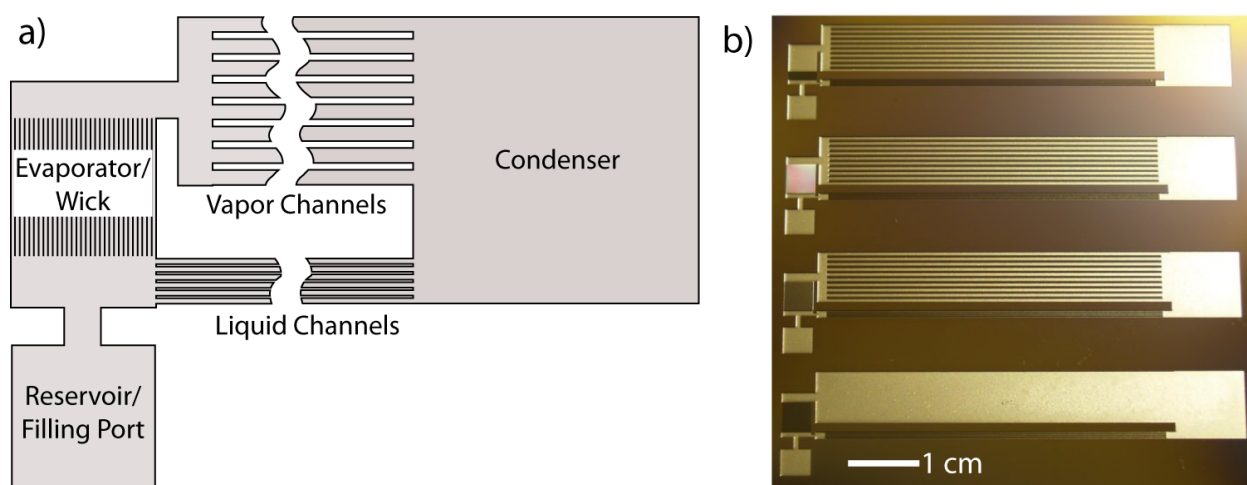


Fig. 4.12 – a) Schematic of the Phase I μ C-LHP layout. **b)** Undiced wafer-level photograph of four Phase I devices with a variety of wick and vapor channel widths.

After the devices (still at 4" wafer level) had been capped with anodically bonded Pyrex, they were diced using an Esec Model 8003 Dicing Saw in the UC Berkeley Microlab (see Fig. 4.13). A glass-cutting blade (part number CX-010-600-060-J from Dicing Technology) was used with a spindle speed of 30,000 RPM and a feed rate of 0.5 mm/s. Holes were then bored through the Pyrex cap into the reservoir region using a 12,000 RPM tabletop drill press fitted with a 1.0 mm diameter electroplated solid diamond drill bit (part number 4ED10 from UKAM Industrial Superhard Tools) and ample water-soluble cooling fluid. (Note: in the first few trials, in order to prevent abraded Pyrex particles from entering the device during boring, I actually bored holes in the Pyrex wafer *prior* to anodic bonding. However, I found that it was *extremely* difficult to bore multiple holes in a Pyrex wafer without fracturing it. Given the expense of Pyrex wafers and the painstaking effort of boring, I decided to instead bond first, then dice, and then lastly bore fill ports into the individual devices. That way, if I did fracture the Pyrex while drilling, I only lost a single device, not an entire wafer.)



Fig. 4.13 – Esec 8003 Dicing Saw.

Initially, I tried to fill the devices with degassed DI water using just a single fill port centered over the reservoir, using a syringe mated to a standard 1.0 mm Upchurch Nanoport assembly. However, I quickly found that it was nearly impossible to force water into the device while simultaneously pushing the air out. Since I wanted the devices completely filled with water for testing, I decided to drill two holes in each device, one centered over the reservoir and one over the condenser. With water entering one hole through the Nanoport and air/water exiting the other hole (and often flushing water back and forth by repeatedly switching the Nanoport from one hole to the other), this method proved far more effective at filling the devices, typically resulting in only a few small air bubbles trapped in the corners of the devices.

4.4.2 – Epoxy-based Sealing and Testing Methods

To seal the two fill holes themselves, I tried three different techniques: Simple Plug, Plug Under Slide, and Ring Under Slide (see Fig. 4.14). With the *Simple Plug* method, I merely covered each hole with a small glob of epoxy and let it cure. The *Plug Under Slide* method added a small (~ 1 cm square) glass slide placed on top of each glob and then compressed between steel parallels with approximately 2.5 lbs of force during curing. *Ring Under Slide* was identical to Plug Under Slide, but rather than a solid glob of epoxy covering the hole, a tiny *ring* of epoxy was put down *around the hole* such that epoxy was not forced into the hole under the pressure of the plates during curing.




Fig. 4.14 – Three epoxy-based dealing techniques shown schematically in cross section with actual samples photographed from above: **a)** Simple Plug; **b)** Plug Under Slide; **c)** Ring Under

Before discussing the results of the three sealing techniques, however, (as well as my reasoning for adopting each of the methods sequentially) it is prudent to first discuss how the filling process was quantified, since ultimately *fluid loss* had to be measured. After boring the fill ports, the devices were flushed with DI water, isopropyl alcohol, and DI water again. This removed the drilling fluid residue and as much of the abraded drilling debris as possible. After fully drying them on a hot plate, the devices were individually weighed using a Denver Instruments precision balance (Model #214), accurate to 0.0001g. (Note: the precision of the balance combined with the total mass of water in a filled device allowed us to calculate mass loss to approximately 1%, more than sufficient to justify satisfaction of the Phase I milestone of 10%

loss.) Each device were then filled with degassed DI water using the methods described above and reweighed. This difference provided us with the mass of the enclosed water, which was recorded. Then (as quickly as possible to minimize evaporation) the two holes were sealed using one of the three techniques described and weighed a third time to determine the mass of the sealing agent and glass cover slides. Since all of the epoxies tested contain some fraction of volatile solvents, the devices were weighed periodically during the curing process (until a steady weight was reached, typically 24-48 hours) to determine the mass of the final cured system.

Once the devices were cured, they were placed into a TPS Environmental Test Chamber (Model TJR, controlled by Partlow MIC 1462, see Fig. 4.15) and kept at 100°C. For the first several days, mass measurements (or simply visual inspections) were performed daily to check for systemic seal failure (which inevitably led to essentially 100% fluid loss). This indicated that the (tremendously variable) manual sealing process had failed for that particular device. If a device “survived” the first week without complete fluid loss, the seal was systemically “good”, and mass measurements were then made approximately weekly. For the Phase I milestone (10% fluid loss over one year at 100°C), DARPA required at least three months of continuous data, from which a year-long estimate could then be extrapolated.



Fig. 4.15 – TPS Tenney Junior Environmental Test Chamber.

4.4.3 – Qualitative Failures and Methodology

For the first set of experiments, I tested five different easily-obtainable consumer epoxies, glues, and sealants (Superglue, JB Weld, *etc.*), using the Simple Plug sealing method for all samples. Again, my reasoning was, why make it more complicated if a simple solution will work? However, I found that some of the sealants failed to prevent leakage even at room temperature, and the others (none of which were designed for elevated temperatures)

systemically failed after a few hours in the 100°C test chamber. The globs either melted or turned into a powder that was easily scraped off. I then purchased a more industrially-respected two-part epoxy called Torr-Seal, designed specifically for vacuum applications. Once again using the Simple Plug technique, I found that the seals themselves were much, much stronger; it was even difficult to scrape off the cured epoxy with a razor blade. However, during curing, cracks appeared in the Pyrex, extending through the full thickness and causing complete fluid loss. Notably, these cracks always originated at the seal sites (see Fig. 4.16) and, moreover, the larger the glob of Torr-Seal, the more extensive the cracking. This led me to believe that either tensile or compressive stresses in the curing epoxy were causing the propagation of microcracks formed during hole boring. I considered etching through-holes in the Pyrex (rather than mechanical boring) to prevent microcracks, but this requires complicated masking. I also attempted to “blunt” any existing microcracks with a quick dip in hydrofluoric acid but quickly abandoned the process after it led to extensive clouding of the Pyrex (Fig. 4.17).




Fig. 4.16 – Microcrack propagation in the Pyrex originating at the Torr-Seal plug.




Fig. 4.17 – HF-dipped Pyrex wafer is heavily clouded. (Note: dotted line superimposed over wafer edge for clarity.)

Based on the observation that bigger globs of Torr-Seal caused more extensive cracking, my next strategy was to try to use less Torr-Seal by implementing the Plug Under Slide technique. In these tests, there were still cracks during curing, but they only appeared from the hole over the condenser; the hole over the reservoir region was almost always crack-free. This strengthened my belief that stresses in the Pyrex were causing propagation of cracks originating at the hole sites. Because the Pyrex “membrane” over the condenser was much larger in area than that of the reservoir, it was also weaker and more prone to fracture under high stress.

At this point we decided that microcrack propagation was probably going to continue to be a problem as long as we bored through the Pyrex mechanically. Why not instead etch fill

ports through the backside of the *silicon* wafer (prior to anodic bonding), rather than through the topside of the Pyrex? Silicon is a stronger material than Pyrex and would perhaps hold up better under the stresses of the curing Torr-Seal. Initially we had avoided this because through-wafer etching takes a long time and is much more costly than boring through Pyrex. However, our next iteration of the Phase I μ C-LHP involved through-wafer thermal isolation trenches, so we would have to perform a through-wafer etch step anyway. Thus my next attempt at sealing used devices with a layout identical to Fig. 4.12, except now with 1 mm square holes etched through the silicon wafer at the corners of the reservoir and condenser. Continuing with the Plug Under Slide technique (but now with glass slide bonded to the backside of the device, on the silicon), we still observed cracks in the Pyrex during curing. Since there were no holes bored in the Pyrex (and thus no microcracks), I posited that perhaps the stress in the Pyrex was due simply to the Torr-Seal expanding during curing. Perhaps the key was to *keep the Torr-Seal out of condenser cavity*, so it can't expand against the Pyrex.

This finally led me to develop the third sealing method, what I called Ring Under Slide. By putting down a small ring of epoxy *around the hole*, rather than a glob of epoxy that gets forced *into the hole*, I could largely limit the stresses put on the Pyrex during curing. This indeed proved to be the case. The devices survived curing without fracture, and I had my first “victory” in this unexpectedly challenging hermeticity study. That victory was short-lived, however. Although the devices survived curing, fluid *gradually* evaporated from the devices once placed in the 100°C test chamber, with the best ones lasting just over 4 days. Moreover, the glass cover slides could be pried off rather easily with a razor blade after the high temperature test. This indicated to me that, although the Ring Under Slide technique was probably the best way to go in terms of seal geometry, Torr-Seal was not the proper material.

4.4.4 – Device Layout and Sealant Changes

Although my work thus far had not produced hermeticity levels anywhere near the DARPA Phase I requirements, I had learned several important lessons and, more importantly, gained valuable experience bonding, boring, and filling devices more efficiently. I also gained much greater proficiency working with epoxies and glass cover slides. Every subsequent trial yielded quicker, cleaner seals. Still, Torr-Seal was not proving adequate. However, in my last set of experiments with Torr-Seal using the Ring Under Slide method, the best results (*i.e.* the devices that retained fluid for the longest time at elevated temperature) were obtained from devices with the least amount of enclosed air (*i.e.* the fewest trapped air bubbles). Because of the small channel diameters and complicated geometry of the Phase I μ C-LHP layout (Fig. 4.12), filling the devices completely with water without a vacuum station was very difficult. Once an air bubble got trapped near the wick or on the wrong side of the transport channels, it was often difficult to dislodge and flush out. This led me to design and fabricate a new set of devices with dimensions and internal fluid volume similar to the Phase I μ C-LHP layout, but with channels designed specifically to be easier to fill and seal. These hermeticity-specific layouts are shown in Fig. 4.18 below. Note that the two fill ports are positioned close together such that one single glass slide can cover them both. This both shortens the evaporation time between filling and sealing and cuts in half the likelihood that a systemically failed seal will lead to complete fluid loss. Also, because both the fill port cavities are small (3 mm square compared to a 10 mm square for the condenser fill port in the Phase I device), any deflection of the Pyrex “membrane” during curing will be reduced, thus decreasing the stress. Lastly, there are no complicated flow

patterns or small channels to trap air bubbles. Because we only needed to demonstrate a successful sealing *method*, it was acceptable to DARPA for us to seal a device that was only *similar* to our actual Phase I device. As it turned out, this small design change made a tremendous difference in both time and results; the new devices were, indeed, *much* easier to fill and seal, particularly the Asymmetric and Serpentine designs.




Fig. 4.18 – Wafer level and enlarged schematics of four hermeticity-specific device layouts.

In addition to these layout changes, I purchased a new single-part epoxy called E6000 that was highly recommended by TAP Plastics. E6000 is manufactured by Eclectic Products and is recommended for just about any sort of adhesive/sealing application on just about any materials. I was initially quite skeptical of a single-part epoxy (as opposed to the two-part base/hardener of products such as Torr-Seal and JB Weld), but it had worked extremely well for the polycarbonate bonding in my Master’s Thesis project (see Appendices).

Filling, sealing (using the Ring Under Slide method), mass measurements, and long-term heating procedures were identical to those described in Sections 4.4.1 and 4.4.2. One small difference was that the parallel plates used to press the glass slides against the devices during curing were compressed using small woodshop C-clamps, rather than simply placing the plates under a 2.5 lb weight. Although I tried to minimize the amount of E6000 forced down into the fill ports by using the Ring Under Slide method, the increased clamping pressure of the C-clamps inevitable led to the holes being clogged. Figure 4.19 is a photograph of two Asymmetric devices sealed with E6000 under a single glass cover slide for each device. The small air bubbles circled in yellow are a result of the epoxy being forced into the hole, thus pushing any air within the “neck” of the hole down into the device itself. The bubbles circled in black, on the other hand, are bubbles that attached to the corners during filling and could not be flushed out.




Fig. 4.19 – Two Asymmetric devices filled and sealed with E6000 under a glass cover slide.

4.4.5 – Quantitative Results

As discussed above, these devices, once filled, sealed, and cured, were placed into a controlled test chamber and maintained at 100°C. Periodic mass measurements were used to determine the amount of fluid lost from the devices. DARPA required at least three months of continuous testing, which could then be extrapolated out to one year for comparison to the milestone. The results of this study are shown in Figs. 4.20 and 4.21. Figure 4.15 shows 90 days of data for five “successful” devices (*i.e.* those that did not fail completely and cause total fluid loss), while Fig 4.21 shows data for three “failed” devices, which did have total fluid loss. In both graphs, the DARPA Phase I milestone of 10% loss is shown in red (and please note the different *y*-axis values). Several points are worth discussing.

First, it is appropriate to address the precision and fidelity of our mass-loss measurements, which formed the bulk of our hermeticity presentation to DARPA. As mentioned above, the Denver Instruments balance used in these measurements has a precision of 0.0001 gram. The individual devices shown in Fig. 4.18 contained between 0.06 and 0.07 grams of fluid when completely filled and sealed. Given the precision of our balance, we could thus resolve fluid losses down to less than 0.2%. For our devices, a 10% fluid loss (as per the DARPA milestone) represents between 0.006 and 0.007 grams. This is 60-70 times larger than the precision of the balance, so we can be very confident in our measurements.




Fig. 4.20 – 90-day fluid loss data for five “successful” E6000-sealed samples. Any reasonable extrapolation to one year would likely satisfy the Phase I milestone of 10% fluid loss.

Second, regarding the three-month (90-day) data shown in Fig. 4.20, it is clear that the fluid loss rate (*i.e.* slope) decreased over time. By the end of the test, the weekly change in mass for many of the devices was below the precision of the balance. Furthermore, although the fluid loss for these five “successful” devices was between 3-6% over three months, the shape of the data does *not* suggest that the year-long fluid loss would simply be four times this amount (*i.e.* 12-24%). The loss rate is clearly nonlinear, and any reasonable extrapolation of the data over one year indicates that we would likely meet the 10% DARPA milestone. This seemingly asymptotic behavior can also be explained theoretically. As discussed in Section 4.1, leakage is a physical phenomenon; it results from fluid molecules being driven through the microscopic and macroscopic structure of the solid by a concentration or pressure gradient. Once a certain amount of fluid was lost, perhaps the driving force (in this case, likely expanding air or vapor in the device) was no longer strong enough to drive more fluid through the resistance of the epoxy.

Next, it must be said that not all of the devices sealed with E6000 met the Phase I milestone. Figure 4.21 shows data for three E600-sealed devices that clearly failed. As I’ve discussed, the painstaking, manual filling and sealing process was highly unrepeatable, and it was not surprising to see some devices fail. What is important to notice, however, is that all three of these devices failed dramatically *within the first day or two*. I never observed a device that was able to maintain fluid levels over 90% for several days or weeks *before* ultimately

failing. Had this occurred, it might have indicated that there was a problem with the E6000 material itself, rather than with my crude, poorly-reproducible sealing technique. Instead, what I observed was the following: *if a seal failed, it failed completely and immediately*. This indicated that failure, when it occurred, was due to sealing *technique*, not sealing *material*. This conclusion is also supported visibly in the seals themselves. Figure 4.22 shows micrographs of two different E6000 seals under 50x magnification. The seal on the left (Fig. 4.22a) was from one of the five “successful” devices (*i.e.* those plotted in Fig. 4.20). Notice how the epoxy around the seal is uniform in color, optically clear, and largely free of voids. On the other hand, the seal on the right (Fig. 4.22b) was from one of the three “failed” devices (those plotted in Fig. 4.21). The epoxy is extensively clouded and riddled with voids, indicating that a proper seal was never formed in the first place. These voids could be due to non-uniform solvent evaporation or perhaps residual water near the port that became mixed with the curing epoxy before ultimately evaporating. Whatever the case, it is not hard to imagine that such a void-riddled seal would sufficiently reduce the driving force necessary to cause complete leakage. Moreover, the epoxy over the hole itself looks dramatically different in the two images. While the successful device appears to have solid, uniform cured epoxy extending down into the hole (the light-colored circle), the hole in the failed device appears to have a full-width void that extends all the way down to the reservoir (the dark-colored circle). While I cannot be certain of the details or the cause simply by examining these pictures, it is clear that the final geometry of the cured “plug” was dramatically different in the two devices.




Fig. 4.21 – 57-day fluid loss data for three “failed” E6000-sealed samples. Note that those devices which failed to meet the 10% fluid loss milestone leaked “catastrophically” within the first two days, *not gradually*.

Finally, although precision mass measurement was the primary (and only truly quantitative) method used to determine system fluid loss, it was not entirely uncorroborated. During testing, it was noted that usually all of the air bubbles in a sealed device would eventually coalesce to form a single bubble. Within the heated test chamber, this bubble usually disappeared entirely (due to the increased gas solubility with elevated temperature). However, during weekly mass measurements (performed at room temperature), it was possible to observe the bubble slowly reform as dissolved gas came out of solution. Moreover, once at room temperature, it was also possible to observe the weekly growth of the bubble, indicating that fluid was lost from the system. This change in bubble size gave us a crude but surprisingly accurate secondary technique to corroborate our mass-loss data.




Fig. 4.22 – 50x micrographs of sealed ports in both **a)** “successful” and **b)** “failed” samples. Contrast the uniform optical clarity and smoothness of a “successful” seal with the frostiness and voids of a “failed” seal.

In conclusion, Fig. 4.20 shows clearly that E6000 *can* work as a Phase I sealant, while Fig. 4.21 seems to indicate that I had simply trouble making the sealing process repeatable. However, DARPA only required us to demonstrate a successful *technique*, and they were sufficiently satisfied with our hermeticity results to award us Phase II funding (particularly when combined with the results from preliminary solder-based sealing methods, discussed in the previous section).

To conclude this chapter, I will again summarize the primary accomplishments and lessons learned regarding port sealing, as well as some general conclusions pertaining to the bonding component of hermetic systems.

- 1) The major structural materials of the μ C-LHP architecture (single crystal silicon and glass) are excellent hermetic materials due to their high density and molecular structure.

- 2) Anodic bonding between glass and silicon layers (particularly when allowing for a wide “bond border” around the functional device area) also forms a hermetic seal.
- 3) Sealing the fill ports will almost certainly be the “weakest link” in the hermeticity game. It is the least automated, the least integrated, and the only “bonding/sealing” step that occurs in the presence of working fluid.
- 4) Epoxy-based sealing methods, the best results of which came with E6000, were sufficient to meet the DARPA Phase I milestones, but cannot be used for later phases.
- 5) Because the “successful” E6000-sealed devices did not explode at elevated temperatures (unlike the solder-sealed devices), it is likely that expanding air (as well as some water) was able to diffuse through the seals during heating. However, the diffusion resistance (due mostly to the width of the seal) must have been sufficiently strong to prevent unacceptable fluid loss (at least for the given concentration difference across the seal).
- 6) Solder seals are very difficult to produce with regularity (at least manually by a novice), but they show excellent promise for effective later-phase DARPA hermeticity milestones.

Chapter 5: The Open-Loop Columnated Microevaporator

At this point, the discussion shifts from essential background material and supporting peripheral studies to the primary focus of this dissertation: the design, fabrication, and testing of an open-loop silicon-based microevaporator for direct integration with the μ C-LHP project. Evaporator design is addressed primarily in the present chapter, while fabrication and testing are reserved for Chapters 6 and 7, respectively. This chapter is divided into three sections. In the first, I briefly describe the overall system architecture and operating principle. In the second, I provide the rationale for employing an open-loop, hollow-columned design in the first place. Specifically, I discuss how an open-loop columnated device allows us to calculate meaningful metrics regarding effective in-plane thermal conductivity and suggest how such experimental data will prove valuable to the greater μ C-LHP effort. Finally, in the third section, I further enumerate the various evaporator designs fabricated and tested in this study and introduce several nondimensional numbers used to characterize their specific columnation schemes.

5.1 Basic Architecture and Operating Principle

The columnated microevaporators employed in this study consist of two silicon dies, each with a variety of features etched into their top surfaces, that are then clamped together (top surface to top surface) to form a two-layer sandwich enclosing a columnated vapor chamber (see Fig. 5.1; dimensioned drawings of the silicon components are shown in Fig. 5.2). Each layer is fabricated on a six-inch 625 μ m thick silicon wafer and then diced to the proper lateral dimensions. The bottom layer, or “floor”, consists of a 10 mm \times 10 mm etched-out evaporator region characterized by an orthogonal array of regularly-spaced raised columns. The columns themselves are hollow, with interior holes etched all the way through the wafer.




Fig. 5.1 – Exploded view and cross-section of open-loop columnated microevaporator architecture.

Leading away from the evaporator region are eight parallel vapor channels, also 10 mm long, and each approximately 1 mm wide. Combining the evaporator and vapor channel regions, the total “device area” of the floor is thus 10 mm × 20 mm, with an additional 2 mm border surrounding the device area for clamping, giving overall dimensions of 14 mm × 24 mm.

The top layer, or “ceiling”, consists of a 10 mm × 10 mm region of uniformly patterned microtextures on an otherwise unpatterned 14 mm × 40 mm die. When properly aligned and clamped, this textured region lies directly overtop the columnated region of the floor and interfaces directly with the tops of the columns. The ceiling layer is 16 mm longer than the floor, thus providing an overhanging region for heat absorption from a ceramic heater that then provides predominantly *in-plane* heat flux to the ceiling of the columnated evaporator region.




Fig. 5.2 – Fully dimensioned drawings of floor and ceiling components.

In terms of basic operation, liquid water is pumped (via servo-driven syringe pump) upwards from a reservoir on the backside of the floor, through the hollow cores of the columns, to the vapor chamber ceiling (see Fig. 5.3). Upon contacting the hydrophilic microtexture of the ceiling, the liquid spreads out over the surface and is vaporized by the externally-provided heat

flux. This newly-formed vapor is ejected from the evaporator via the parallel vapor channels, terminating and condensing at a fixed-temperature cold sink at the end of the channels.

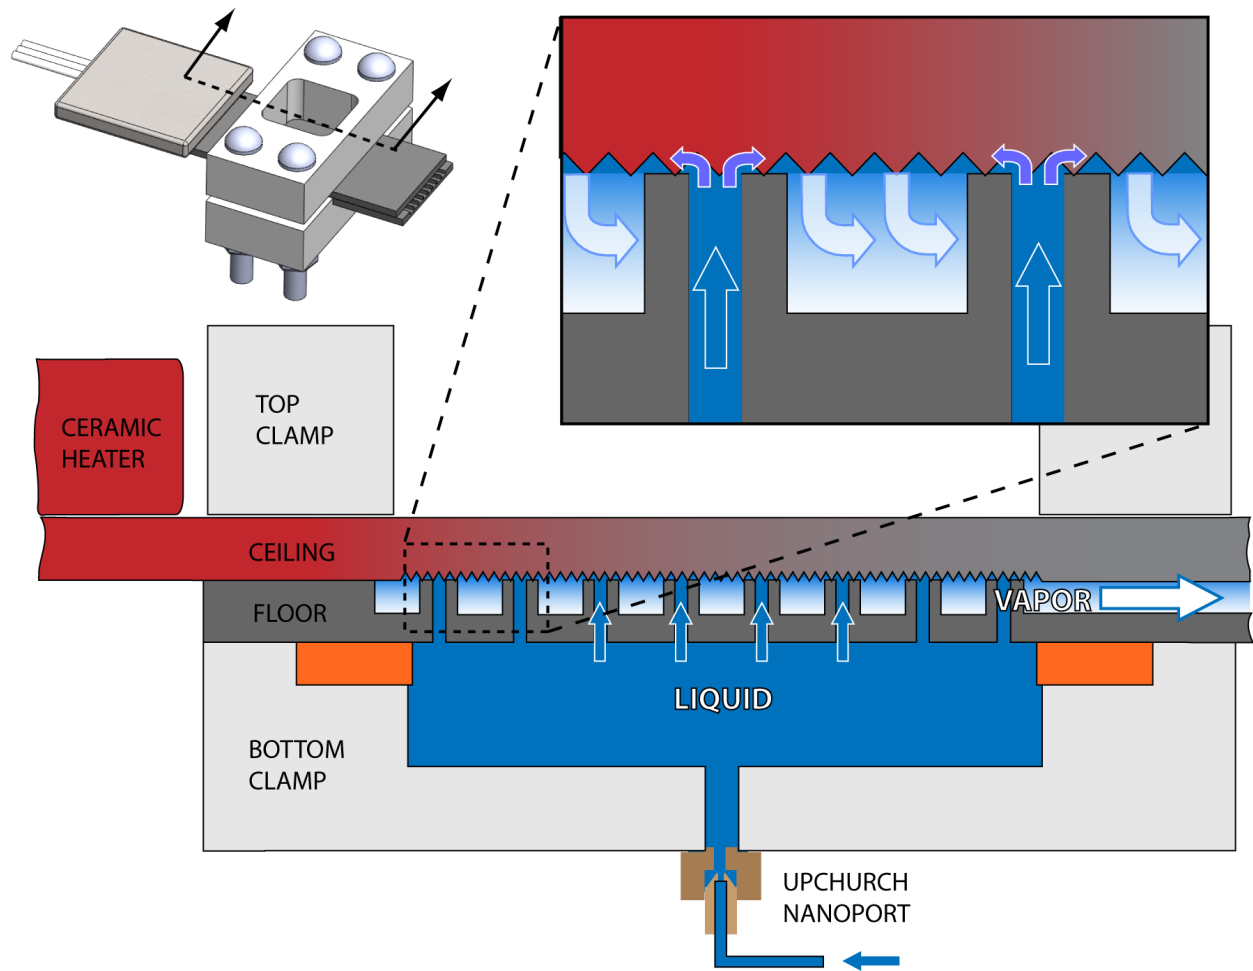


Fig. 5.3 – Cross-sectional schematic of microevaporator illustrating basic operation.

The theoretical motivation is as follows. Under steady state conditions, the applied heat flux is dissipated along two paths, shown schematically in Fig. 5.4a. The first path, solid conduction, is governed (to first order) by the thermal conductivity k of single crystal silicon and the in-plane cross-sectional area of the device. The second path, convection with phase change, is governed (again, to first order) by the effective heat transfer coefficient h at the vaporization surface and the area of the microtextured ceiling. Therefore, the larger h becomes compared to k , the greater the vaporization rate and the greater the heat transfer to the convected vapor. Finally, the greater the heat transfer to the convected vapor, the smaller the heat transfer due to conduction, and thus the lower the temperature gradient across the evaporator—which ultimately signifies a higher effective in-plane thermal conductivity. For example, in Fig. 5.4b, a low convection coefficient (and thus low vaporization rate) means that solid conduction through the silicon will dominate and will necessitate a large in-plane temperature gradient across the evaporator. A large temperature gradient implies a low effective thermal conductivity. At the other extreme, in Fig. 5.4c, a high convection coefficient (and thus high vaporization rate) means that solid conduction through the silicon will be minimal, and the necessary temperature gradient

will be small. This, in turn, implies a very large effective thermal conductivity. Although it is often quite difficult to predict the convective heat transfer coefficient *a priori*, a simple one-dimensional model can be constructed to quantify the effective in-plane thermal conductivity of the device *as a function of h* at the vaporization surface (see next section). Although *h* won't be known exactly, there is an accepted range of values based on experimental data, and one can examine how the overall effective in-plane thermal conductivity varies with *h*.




Fig. 5.4 – a) Cross-sectional schematic illustrating the two transport paths for evaporator heat; **b)** Poor convection from the evaporator surface forces increased solid conduction, larger temperature gradients, and thus lower effective thermal conductivity; **c)** Strong convection from the evaporator surface causes decreased solid conduction, lower temperature gradients, and higher effective thermal conductivity.

5.2 Analogous Heat Transfer and Performance Analysis

Having described the basic architecture, operating principle, and *theoretical* motivation, it is now sensible to provide the *practical* rationale for fabricating and testing such a device. After all, the μ C-LHP in its final form is a hermetically-sealed, closed-loop system providing fluidic pumping via capillary forces—not a mechanically-clamped, open-loop system driven by external syringe pump. However, there are a number of reasons why a more realistic architecture was not chosen for this particular design stage. Some of the modifications were deliberate; others were clearly not ideal but instead necessitated by fabrication issues. For instance, it may seem strange to test a hollow-columned design when the μ C-LHP design calls for solid (though porous) columns made of through-wafer coherent porous silicon (CPS). Unfortunately, at the time of this study, the CPS etching process for the μ C-LHP project had yet to produce any samples suitable for device fabrication. However, because the Columnated Vapor Chamber (CVC) concept is the most novel and unique feature of the μ C-LHP, we believed it absolutely essential to begin preliminary testing of a columnated device—even if the structures could not be fabricated in CPS. Furthermore, we knew that suitable CPS *had* been successfully produced and implemented in functional microscale loop heat pipes [5.1, 5.2], so we did not see this as a permanent obstacle to the final design process. Of course, without a porous wick to provide capillary pumping, we were forced to rely on a syringe pump for *mass* transport; nevertheless, from a *heat* transport perspective, the system was designed to be structurally and functionally analogous to the closed-loop, columnated μ C-LHP design.

The next major modification was the decision to mechanically clamp the floor and ceiling layers together for testing, rather than permanently and hermetically bond them. There were a number of reasons for this. First, the KS Bonder tool in the Berkeley Microlab was found to be extremely unreliable during the hermeticity studies described in Chapter 4, often producing a successful bond less than 50% of the time. Moreover, “unreliable” in this context does not mean simply that the wafers were not properly bonded, but rather that the wafers were completely destroyed. Obviously, for the hermeticity studies, this processing step could not be avoided because it was the central principle being tested. Furthermore, the devices being bonded were much simpler to fabricate (one lithography step, one etch), so the occasional shattered wafer was not an overwhelming setback. However, for the much more fabrication-intensive columnated evaporators tested in this study, a 50% failure rate was simply not acceptable. Second, permanent bonding of floor to ceiling dies would greatly limit the number of different column-microtexture combinations on each wafer. Essentially, each floor could be paired to just one ceiling. On the other hand, by allowing interchangeable floors and ceilings to be tested in a reversible clamp setup, the number of testable configurations was greatly increased. In addition, because the devices are completely interchangeable, duplicates fabricated on the same wafer could serve as backups if devices were damaged. A third reason for non-permanent bonding was to facilitate continued access to the interior surfaces of the vapor chamber. This greatly reduced the required purity of the working fluid used during testing and made cleaning and defouling much easier. Also, it left open the possibility of experimenting with additional surface treatments and hydrophilic coatings. Finally, by employing an open-loop design, I could still test many variations of the CVC concept without continually facing the extremely challenging task of fluidic charging and fill-port sealing under vacuum.

What the open-loop columnated microevaporator ultimately provides, then, is a highly flexible platform to test the CVC and IROES concepts, which is both structurally and functionally analogous to the columnated wick structure of the μ C-LHP. Although the *mechanism* of fluid flow is different, the overall heat transfer phenomena are quite similar. In both the open-loop test platform and the final μ C-LHP design, liquid flows up through regularly-spaced columns to the microtextured ceiling, where it is spreads into a thin film due to capillary forces (see Fig. 5.5). Some of the supplied thermal energy will be conducted laterally through the solid *due to an imposed temperature gradient*, while the rest will be absorbed by the latent heat of the working fluid during vaporization and convected more or less *isothermally* away from the evaporator. Consequently, the greater the fraction of thermal energy transferred out-of-plane to the working fluid, the lower the temperature gradient across the evaporator. Furthermore, for a fixed heat input, *a lower temperature gradient across the evaporator necessarily implies a higher effective thermal conductivity*. Broadly speaking then, regardless of the fluid pumping mechanism (whether via capillarity or external pressure gradient), the effective in-plane thermal conductivity of the evaporator will be largely a function of the liquid vaporization rate.




Fig. 5.5 – Comparison of the μ C-LHP evaporator architecture with that of the open-loop microevaporators of this study. Sectional views below each illustrate the analogous columnar structure and fluid flow.

Fig. 5.6 below shows a schematic cross-section of the ceiling of the device and its relation to the entire two-layer microevaporator. The total device length is L , while the evaporator length is L_E . A uniform heat flux q_0'' is applied to the left end over the device to simulate the in-plane flux provide by the ceramic heater on the overhanging segment of the

ceiling. Since the thickness t is much less than the length L or width w , it is reasonable to assume isothermal cross sections (*i.e.* a classic extended fin). Furthermore, it is assumed that the vapor channels are adiabatic (a common idealization in heat pipe analysis) and that the natural convection from the top exterior surface is negligible compared to the phase change convection from the interior surface. Finally, a fixed temperature boundary condition is applied at the right end of the device to simulate the cold sink at T_0 . The most important part of the model is the vaporization surface. Here, a convection boundary condition is imposed, with a variable heat transfer coefficient h . Depending on the specific column density (see Eqs. 5.21 and 5.22 below for relevant metrics), the actual vaporization area will likely be much larger than L_E^2 (due to the topography of the microtextured surface), but in this model it is conservatively assumed to be flat. The governing equations and boundary conditions are summarized below, starting from the approximation of the ceiling layer as a two-part extended fin with piecewise boundary conditions and uniform junction temperature and flux (see [5.3] for derivation).




Fig. 5.6 – Geometry and boundary conditions of a simple one-dimensional model to predict the effective in-plane thermal conductivity of the evaporator as a function of the convection heat transfer coefficient h .

Domain I – Convective Flux from Bottom Surface of Ceiling (from $x = 0$ to $x = L_E$)

We begin with a one-dimensional heat equation for an extended surface,

$$\frac{d^2 T}{dx^2} - \left(\frac{1}{A_c} \frac{h}{k} \frac{dA_s}{dx} \right) (T - T_\infty) = 0, \quad (5.1)$$

where A_c is the cross-sectional area and A_s is the lateral surface area. For the prismatic ceiling layer, A_c is simply the thickness t times the width w , and the lateral surface area derivative with respect to x is simply the width w (not $2w$ because only convection from the bottom surface is considered). Simplifying, we have

$$\frac{d^2 T}{dx^2} - \left(\frac{1}{tw} \frac{h}{k} w \right) (T - T_\infty) = 0. \quad (5.2)$$

Defining $m^2 \equiv \frac{h}{tk}$ and $\theta_l \equiv T - T_\infty$, the governing equation (5.2) becomes

$$\frac{d^2 \theta_I}{dx^2} - m^2 \theta_I = 0. \quad (5.3)$$

Eq. 5.2 takes on the boundary conditions:

$$-k \frac{d\theta_I}{dx} \Big|_{x=0} = q_0'' \quad (5.4)$$

$$-k \frac{d\theta_I}{dx} \Big|_{x=L_E} = q_J'' \quad (5.5)$$

$$\theta_I \Big|_{x=L_E} + T_\infty = T_J \quad (5.6)$$

where the subscript “J” refers to the junction between two domains (*i.e.* at $x = L_E$). The general solution to Eq. 5.3 is

$$\theta_I(x) = A \exp(mx) + B \exp(-mx). \quad (5.7)$$

Enforcing boundary conditions 5.4 – 5.6 yields:

$$-k \frac{d\theta_I}{dx} \Big|_{x=0} = -k mA + k m B = q_0'' \quad (5.8)$$

$$-k \frac{d\theta_I}{dx} \Big|_{x=L_E} = -k mA \exp(mL_E) + k m B \exp(-mL_E) = q_J'' \quad (5.9)$$

$$\theta_I \Big|_{x=L_E} + T_\infty = A \exp(mL_E) + B \exp(-mL_E) + T_\infty = T_J \quad (5.10)$$

Domain II – Both Top and Bottom Surfaces Insulated (from $x = L_E$ to $x = L$)

Again starting with a one-dimensional heat equation for an extended surface,

$$\frac{d^2 T}{dx^2} = 0, \quad (5.11)$$

and defining $\theta_{II} \equiv T - T_0$, the governing equation becomes

$$\frac{d^2 \theta_{II}}{dx^2} = 0. \quad (5.12)$$

Eq. 5.12 takes on the boundary conditions:

$$\theta_{II} \Big|_{x=L} = 0 \quad (5.13)$$

$$-k \frac{d\theta_{II}}{dx} \Big|_{x=L_E} = q_J'' \quad (5.14)$$

$$\theta_{II} \Big|_{x=L_E} + T_0 = T_J . \quad (5.15)$$

The general solution to Eq. 5.12 is

$$\theta_{II}(x) = Cx + D . \quad (5.16)$$

Enforcing boundary conditions 5.13 – 5.15 yields:

$$\theta_{II} \Big|_{x=L} = CL + D = 0 \quad (5.17)$$

$$-k \frac{d\theta_{II}}{dx} \Big|_{x=L_E} = -kC = q_J'' \quad (5.18)$$

$$\theta_{II} \Big|_{x=L_E} + T_0 = CL_E + D + T_0 = T_J . \quad (5.19)$$

The six linear equations 5.8 – 5.10 and 5.17 – 5.19 can then be solved for the six constants A , B , C , D , q_J'' , and T_J .

Following this one-dimensional model, the simulated temperature distribution across the device has been plotted in Fig. 5.7 for fixed heat flux q_0'' and the device dimensions (Note: this particular simulation was done for an 8 cm device, but the basic principal and resulting trends hold for other geometries). It was assumed that the vapor was at saturation and the cold sink was maintained at 0°C. The various curves reflect varying convective heat transfer coefficients. As expected, when h is very low (*i.e.* no vaporization off the bottom surface), the temperature distribution is approximately linear. This reflects the limiting case of purely solid conduction through the silicon. On the other hand, as h is increased (*i.e.* as more liquid is vaporized), the temperature gradient across the evaporator becomes smaller and smaller. In the limit that h becomes infinite, the gradient approaches zero (*i.e.* the evaporator is isothermal across its length). This is the idealization of the DARPA Thermal Ground Plane, with heat being dissipated laterally across an isothermal plane. Note also that the temperature gradient across the vapor channels is always linear, reflecting the fact that only conduction is at work, since the channels were assumed to be adiabatic

From the results of Fig. 5.7, one can easily calculate the temperature gradient itself (say, at the center of the evaporator) as a function of h . Figure 5.8 shows such a plot, for convective heat transfer coefficients ranging from 10-100,000 W/m²·K. Consistent with the previous figure, the temperature gradient rapidly approaches zero (*i.e.* the region becomes isothermal) as the convective heat transfer coefficient approaches 10,000 W/m²·K. Finally, this temperature gradient can be combined with the prescribed input heat flux and device geometry to estimate the effective in-plane thermal conductivity at the evaporator midpoint as a function of h (see Fig. 5.9). Although h is generally unknown (see discussion above), one can still predict the likely operating range of the device by matching the effective thermal conductivities associated with the range of reasonable h values for convection with phase change. As is readily seen in Fig. 5.9, even moderate convection heat transfer coefficients ($h = 500$ -5,000 W/m²·K) put us well within reach of even the most optimistic DARPA conductivity goals. For comparison, Incropera and DeWitt suggest average h values between 2,500-100,000 W/m²·K for convection with boiling [5.4].




Fig. 5.7 – Theoretical temperature profiles as a function of the average heat transfer coefficient h from the underside of the evaporator ceiling, for the one-dimensional evaporator model shown in Fig. 5.5.




Fig. 5.8 – Theoretical temperature gradient at the evaporator midpoint as a function of the average heat transfer coefficient h on the underside of the evaporator ceiling.




Fig. 5.9 – Theoretical effective thermal conductivity of evaporator model shown in Fig. 5.5 as a function of the convective heat transfer coefficient h .

With Fig. 5.9 in mind, it should finally be clear why my ultimate goal is to maximize the average heat transfer coefficient on the underside of the ceiling. Stated simply: *increased convective heat transfer from the underside of the evaporator ceiling directly and dramatically increases the effective in-plane thermal conductivity of the evaporator.* Furthermore, given the superior heat and mass transfer characteristics of the interline region, it should also be clear why I want to vary the columnation and microtexture schemes in order to increase the percentage of the ceiling area with film thicknesses that correspond to the interline meniscus region (see Fig. 5.10). On one hand, if the columns are too close together (for a particular ceiling microtexture), too much of the intercolumnar film area will be too thick for efficient conduction. On the other hand, if the columns are too far apart, too much of the extended meniscus will lie in the adsorbed region, where dispersion forces dominate and vaporization is completely prevented. Additionally, for a device fabricated from coherent porous silicon, the columnation scheme will partially determine the mass flow distribution between ceiling and floor (see Fig. 1.8).

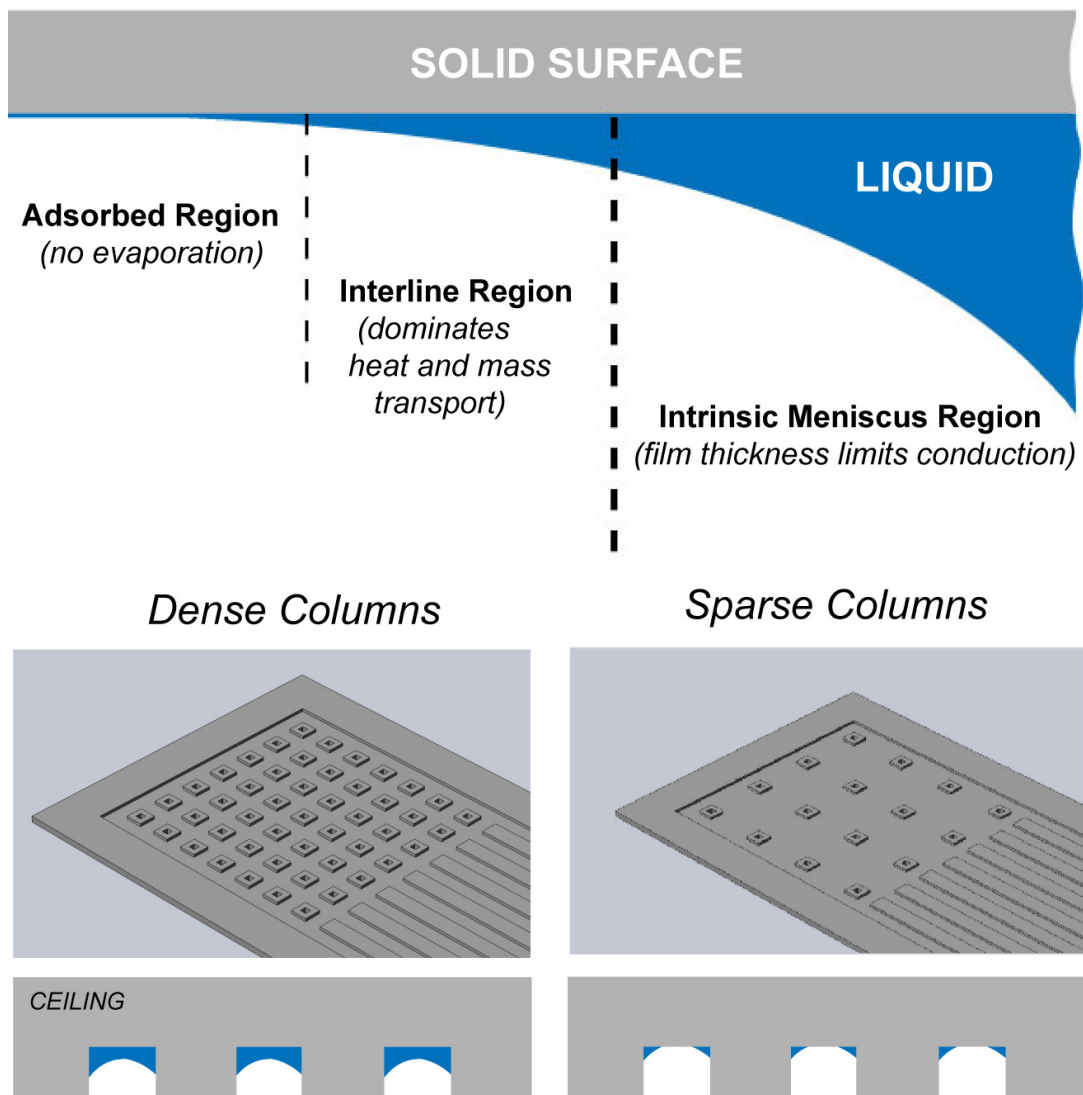


Fig. 5.10 – Schematic illustrating the effect of columnation scheme on the average intercolumnar film thickness.

As discussed earlier, it is extremely difficult to experimentally observe and quantify interline region vaporization, even when using a highly idealized experimental setup designed specifically to do so. This is primarily due to the extremely short length of the interline region. Dhavaleswarapu *et al.* [5.5] suggest a length of only 50 μm for the entire “micro region” (see Fig. 1.10), while the true interline region is only hundreds of *nanometers* in length. Park and Lee [5.6] are in agreement, suggesting that the interline region accounts for approximately 0.5% of the entire extended meniscus. Obviously, attaining this kind of spatial resolution for temperature measurements is extraordinarily challenging. Liquid crystal thermography was once the most popular technique (see [5.7, 5.8, 5.9]), but recent advances in “affordable” IR cameras have allowed spatial resolution down to less than 10 μm [5.5].

Needless to say, quantifying interline region vaporization in an *actual device* is even more daunting, especially given that the IR camera used in this study has a maximum spatial resolution of 120 μm per pixel. However, because the transport enhancement in the interline region is so dramatic (up to 15,000 $\text{W}/\text{m}^2\cdot\text{K}$; see [5.6]), it is hoped that even a slight (perhaps even un-measurable) increase in interline area will still provide measureable effects on the overall performance (*i.e.* effective in-plane conductivity) of the evaporator. As emphasized repeatedly, this study revolves around functional device fabrication. Whether the interline vaporization can truly be quantified or understood in terms of columnation scheme is of secondary importance to the design of a high-performance device. Thus, what I realistically hope to develop during the experimental process is a basic understanding of general performance *trends* as they relate to columnation scheme, heat flux input, and mass flow rate. Such guidelines and upper/lower performance limits should prove extremely valuable to the further development of the μC -LHP program.

5.3 Device Layouts and Design of Experiments

In this section, I explicitly describe the layouts of the numerous floor and ceiling components that were fabricated and tested in this study. Because this represents the very earliest attempt to test the CVC concept, I had little to no guidance with which to constrain initial designs. Should a 10 mm square evaporator have four columns or four hundred? Clearly this is much different than determining the optimum dimensions of a cantilever beam. Any attempt to accurately model a system with such complicated fluid dynamics and phase-change heat transfer would constitute a doctoral thesis in its own right, while any simple, assumption-heavy model would likely prove no more useful than “engineer’s intuition”. As a result, the various design parameters in this study span a large range of values with rather coarse incremental divisions.

5.3.1 – Mask Layout Considerations

Moreover, whenever microfabrication is involved, there are always two diametrically opposed design philosophies at play when it comes to mask layout. On the one hand, it is usually advantageous to have as many *different* designs as possible on a single mask (*i.e.* fine incremental divisions of a particular parameter) in order to most accurately determine the functional dependence of each design variable. On the other hand, *multiple instances* of the *same* design on a mask provide insurance in the event that devices are damaged during fabrication or testing, but this obviously requires coarser incremental divisions. Had I permanently bonded the floor and ceiling wafers prior to dicing, I could have produced only as many different floor-ceiling combinations as could fit on one 6” wafer (approximately 20).

Moreover, since the floor dies are significantly smaller than the ceiling dies, permanent bonding prior to dicing would have reduced the absolute number of floor dies per wafer and wasted a large amount of the floor wafer area. By keeping the floor and ceiling layers interchangeable, however, I not only increased the absolute number of floor dies per wafer, but also greatly increased the number of testable device combinations obtainable from a single mask set. This greater freedom also allowed me to include multiple instances of many of the floor and ceiling design without requiring even coarser incremental divisions. That being said, even with these improvements, an optimal layout is not likely to be found among the fabricated designs. However, given the preliminary nature of this design stage, true optimization in the strictest sense was never a realistic goal. Instead, the practical objective was to establish the functional dependence of the design variables (even if only qualitatively) and hopefully determine some general trends in evaporator performance to guide later, more refined designs. Wafer-level schematics illustrating die layout can be seen in Figs. 5.11 and 5.12 below.




Fig. 5.11 – Wafer-level mask layout for device floor fabrication. The layout contains 36 standard-sized floor dies (14 mm wide by 24 mm long) utilizing 12 different columnation schemes, along with one elongated die with 20 mm channels. The column patterns are too dense to be seen at this magnification, but the cyan regions indicate the active device areas. Zoomed-in views of individual die designs can be found in Chapter 6.

5.3.2 – Floor Designs

For the floor components, the primary design variables were column size and intercolumnar spacing. As discussed above, I based the initial range (particularly the upper limit) largely on intuition (*e.g.* one gigantic column would render the whole columnated wick concept moot), while lower limits were determined mostly by fabrication limitations. Since the interior column holes had to be etched all the way through the wafer, the minimum size was limited by the maximum aspect ratio of the DRIE process. In the end, I chose four exterior column widths d : 200, 300, 500, and 900 μm . In all designs, the column walls are 50 μm thick, thus the interior through-holes are 100, 200, 400, and 800 μm wide. For intercolumnar spacing, I chose three values (center to center pitch p): 1250, 2500, and 5000 μm (see Fig. 5.13). Combining all four column sizes with all three pitches, the mask layout contained twelve different floor designs, shown schematically in Fig. 5.14. Since the wafer layout shown in Fig. 5.11 above allows for thirty-six standard dies, three instances of each design were included. Although it was not explicitly intended as a design variable, it should be noted that, column height (or equivalently, vapor chamber depth) will vary somewhat from one device to the other because of DRIE etch rate nonuniformity. The *bull's eye effect*, for example (discussed more fully in Chapter 6), causes devices at the center of the wafer to be etched more slowly, resulting in shallower channels.




Fig. 5.13 – Columnation parameter definitions and fixed dimensions.

In addition to specifying the column dimensions and pitch, I also assigned two dimensionless parameters to each layout to provide a single variable characterization of the general columnation scheme. The first, called *Columnar Area Percentage* (CAP), is defined:

$$CAP \equiv \frac{\text{Total Footprint of Columns}}{\text{Total Evaporator Area}} \times 100\% \quad (5.21)$$

For the twelve layouts, CAP values ranged 0.16% to 51.84%. The second parameter, called *Column-to-Floor Ratio* (CFR) is defined:

$$CFR \equiv \frac{\text{Total Footprint of Columns}}{\text{Total Footprint of Floor}} \quad (5.22)$$

CFR values for the layouts ranged from 0.0016 to 1.0764. Note that some layouts which have very different column size and pitch still have similar CAP and/or CFR values. I thought that perhaps this device characterization might prove useful in analyzing performance trends. Finally, I also included one die with 20 mm vapor channels (as opposed to the standard 10 mm) to see what effect channel length had on performance.




Fig. 5.14 – Schematic of evaporator floor columnation schemes. See Table 5.1 below for details.

Table 5.1 – Summary of Floor Columnation Schemes.


Design Number*	Column Width: d (μm)	Column Pitch: p (μm)	CAP	CFR
1	200	1250	2.56%	0.0263
2	300	1250	5.76%	0.0611
3	500	1250	16.0%	0.1905
4	900	1250	51.84%	1.0764
5	200	2500	0.64%	0.0064
6	300	2500	1.44%	0.0146
7	500	2500	4.0%	0.0417
8	900	2500	12.96%	0.1489
9	200	5000	0.16%	0.0016
10	300	5000	0.36%	0.0036
11	500	5000	1.0%	0.0101
12	900	5000	3.24%	0.0335

*See Fig. 5.14 above for corresponding schematic.

5.3.3 – Ceiling Designs


For the ceiling components, the two primary design variables were channel size and spacing on three different microtexture geometries: parallel channels, orthogonal channels, and oblique channels (see Fig. 5.15). Once again, the lower limits for feature size were set by fabrication limitations, while the upper limits were based on the microtexture interface with the tops of the columns. For example, it wouldn't make sense to have channels larger than 50 μm , since that is the thickness of the column walls. In the end, I chose four channel widths: 2, 4, 8, and 16 μm . The range of channel spacing was chosen based largely on practical experience with macroscale heat exchangers and common sense (*e.g.* if the channels are spaced too far apart, one ceases to have a microtexture). Combining all four channel sizes and spacings with all three geometries, the mask layout contained twenty different ceiling designs, shown in tables at the bottom of Fig. 5.15. Since the wafer layout shown in Fig. 5.12 above allows for twenty-one standard dies (as well as a several odd-sized ones), I was able to include multiple instances of several intuitively “promising” design. Also, an elongated ceiling die was included to accommodate the long-channeled floor die. Finally, as mentioned in the previous section (and discussed more thoroughly in Chapter 6), channel depth will also vary from one device to another due to etch rate nonuniformity.

Parallel Microchannels




Channel Width (μm)	Channel Spacing (μm)
2	2
2	4
4	4
8	4
8	8
16	4
16	8
16	16

Orthogonal Microchannels



Channel Width (μm)	Channel Spacing (μm)
2	2
2	4
4	4
8	4
8	8
16	4
16	8
16	16

Oblique Microchannels



Channel Width (μm)	Channel Spacing (μm)
2	2
4	4
8	8
16	16
All 30° offset	

Fig. 5.15 – Geometry and channel parameters for all twenty ceiling microtexture designs.

Chapter 6: Columnated Microevaporator Fabrication

This chapter describes in moderate detail the various processes used to fabricate the primary structural components of the two-layer open-loop microevaporator described in Chapter 5. These components are divided between two silicon layers, which hereafter will be referred to simply as the *Floor* and *Ceiling* of the microevaporator, for obvious reasons (see Fig. 6.1). This chapter is divided into three sections; the first section details the final fabrication of the Floor, while the second section describes the final fabrication of the Ceiling. Finally, a third section is devoted solely to describing some of the technical difficulties, mishaps, calibration steps, and process modifications that accompanied some of the processes. In my description of the various fabrication processes (deposition, etching, *etc.*), I will assume the reader has a basic understanding of microfabrication techniques and will not attempt to discuss the physics or chemistry of the reactions. Detail regarding the underlying science is readily available in various introductory texts ([6.1, 6.2], for example), while further details regarding the performance or configuration of the particular tools used in this study can be found in the online manuals of the Berkeley Microfabrication Laboratory and the Marvell Nanofabrication Laboratory [6.3, 6.4].




Fig. 6.1 – Nomenclature used for the two structural layers of the microevaporator.

This chapter will also not be a “cookbook” that details the process flow to such an extent that any researcher could replicate the fabrication (*e.g.* lamp intensity, exposure time, spin rate, RF power). There are three reasons for this. First, as anyone who has ever worked in a university microfabrication facility can verify, a specific recipe that worked a month ago is unlikely to work today. There are simply too many parameters that cannot be controlled and too much variation that cannot be accounted for. Furthermore, process variables for one wafer may be vastly different from those of a second wafer, even though the final result (and thus the overall process flow) is identical. Secondly, I have included as Appendix C a step-by-step process flow (in a more concise “recipe-like” format) that further details the processing steps *when appropriate*. I cannot emphasize the words “when appropriate” enough. For example, it would be superfluous to inform the reader that I first etched for 10 minutes, *then* performed a

profilometry measurement to calculate the etch rate, *and then* continued the etch to a depth of 120 microns. Such characterization steps are simply understood as requisite to microfabrication, and the details and conditions *that particular day* are neither helpful to future research nor indicative of the way microfabrication actually works in practice. In the aforementioned example, for instance, I'll simply state that I etched 120 microns. Lastly, in several instances, I describe only the *final version* of a particular processing step within Sections 6.1 and 6.2, while setting aside alternative earlier iterations and processing techniques for Section 6.3. This helps the primary fabrication discussion maintain a more linear, logical, and uninterrupted flow.

6.1 Microevaporator Floor Fabrication

Floor fabrication was performed on 6" p-type prime silicon wafers with a nominal thickness of 625 μm . The process involved two lithography steps, one using a patterned thin-film oxide hard mask, while the other a multilayer photoresist soft mask. Schematic cross-sections of all processing steps can be seen in Figs. 6.2, 6.4, 6.8, and 6.11. After pre-furnace cleaning (Piranha solution then hydrofluoric acid), the wafers were placed in a Tystar Low Pressure Chemical Vapor Deposition (LPCVD) furnace, and approximately 2 μm of undoped low-temperature oxide (LTO) was deposited at 450°C. After depositing a monolayer of hexamethyl disilazane (HMDS) to improve photoresist adhesion, the wafers were coated with 1.3 μm of G-line positive photoresist that was then soft-baked for 60 seconds at 90°C. Using Mask #1 (see Fig. 6.3a) in a Karl Suss MA6 Mask Aligner, the photoresist was then exposed to define the channels and columns. A post-exposure bake for 60 seconds at 120°C was followed by standard G-line development, thus exposing the oxide for hard mask definition. Finally the photoresist mask was UV hard-baked using a Fusion M150PC Photostabilizer System.




Fig. 6.2 – Schematic wafer cross-sections for Floor process flow (Steps 1 – 4).

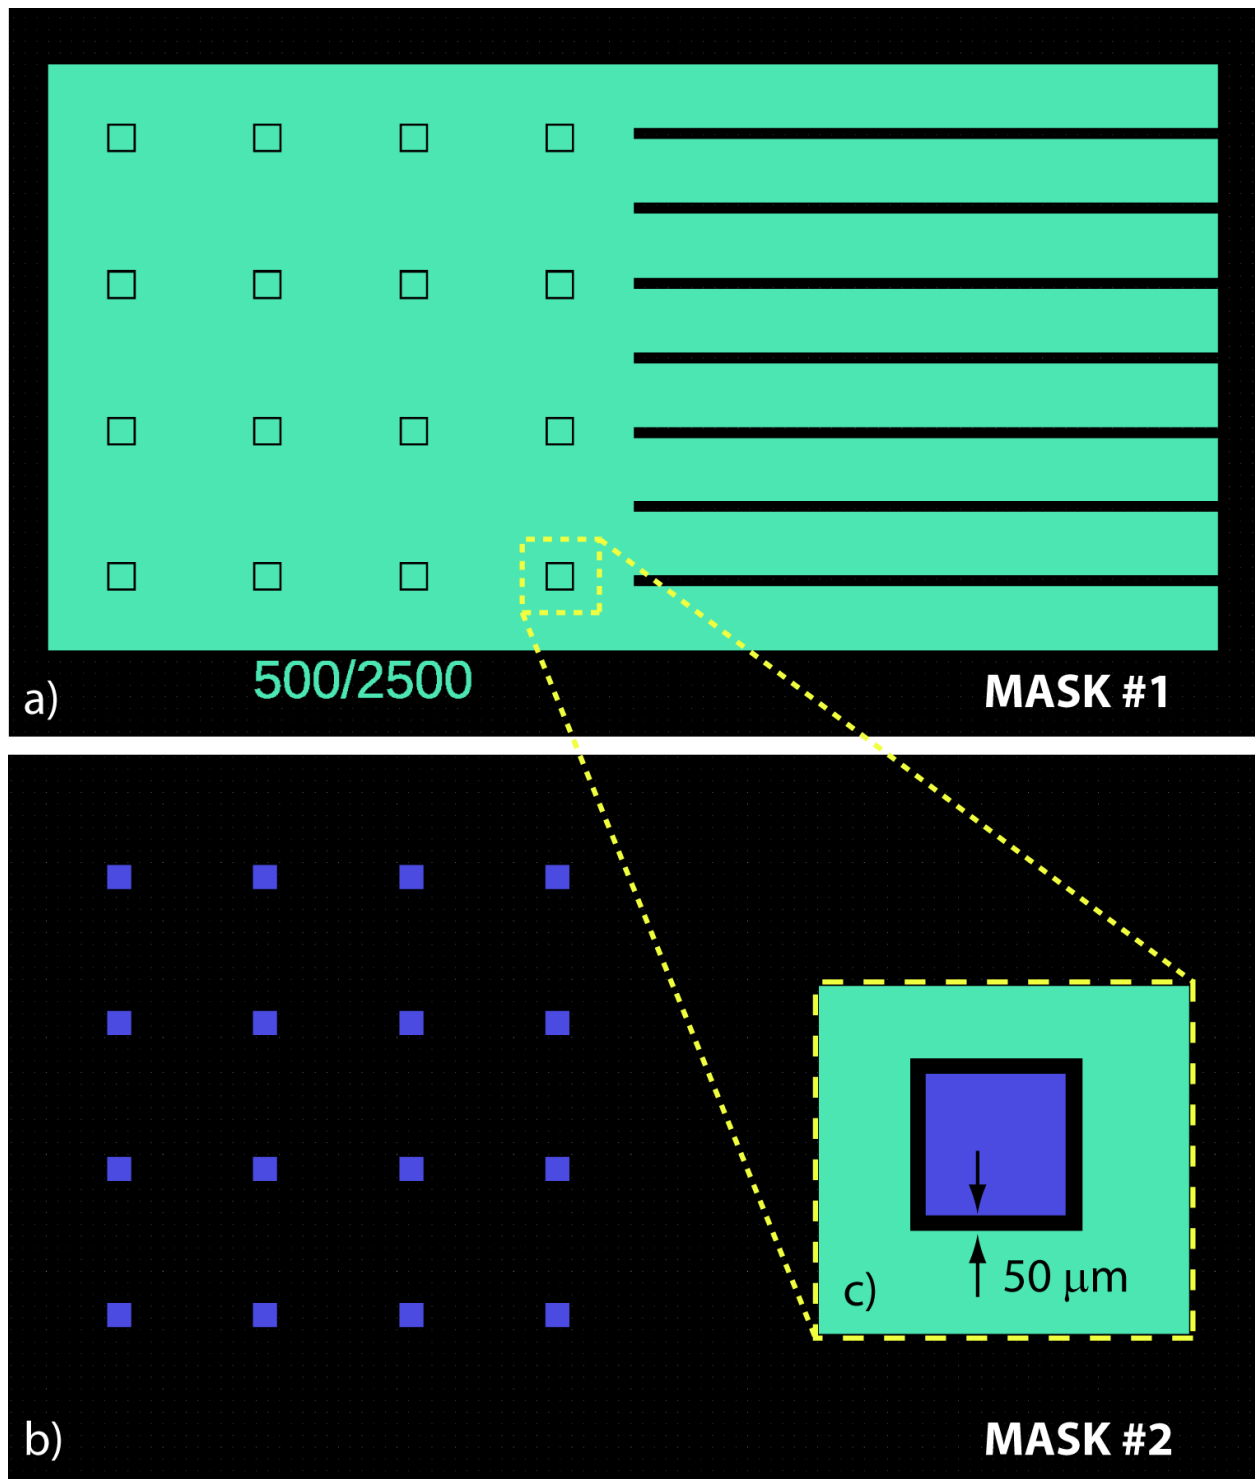


Fig. 6.3 – Characteristic dark field photomasks shown at the device-level. Colored areas are clear on mask and expose/remove positive resist; black areas on mask are chrome and shield resist: **a)** Mask #1 defines the channels and column walls; **b)** Mask #2 defines the through holes; **c)** Zoomed-in view of a single column detailing the hole-column alignment. The black rim between the colored regions demarks the column walls, which are 50 μm thick for all layouts. (Note: the label “500/2500” identifies this device as having columns 500 μm wide, spaced 2500 μm apart, center to center.)

The exposed oxide was then etched down to the bare silicon using CF_4 and CHF_3 in a Centura MxP+ plasma etcher. This patterns the oxide hard mask, which in turn defines both the channels and the interior and exterior walls of the columns. After stripping the photomask in a PRS-3000 bath heated to 80°C , the wafers were again cleaned in Piranha and coated with HMDS in preparation for the second lithography. This lithography step required nonstandard processing because the column holes had to be etched all the way through the wafer (or at least nearly so). This requires a very robust etch mask. Since the STS etch selectivity between silicon and photoresist is about 50:1, I chose to use 5 layers of $2\text{ }\mu\text{m}$ G-line positive photoresist (each soft-baked for 90 seconds at 90°C) for a total mask thickness of approximately $11\text{ }\mu\text{m}$. Using Mask #2 (see Fig. 6.3b) in a Karl Suss MA6 Mask Aligner, the photoresist was exposed to define the through holes. Note that the through holes are the same size as (and aligned directly over top of) the inner column walls defined by Mask #1 (see Fig. 6.3c). Given the extreme thickness of the resist, very long exposure times were needed to expose to full depth, and the post-exposure bake was eliminated to avoid thermally stressing the thick layer. Two standard passes of G-line developer were required for full development. Finally the thick photoresist mask was hard-baked for five hours at 120°C because a standard UV bake would not penetrate through the entire mask. Micrographs of the characteristic device features at this point in the processing can be seen in Fig. 6.5.




Fig. 6.4 – Schematic wafer cross-sections for Floor process flow (Steps 5 – 8).




Fig. 6.5 – Micrographs of the characteristic features of a device after hard-baking the thick resist soft mask (Step 8): **a)** 200 μm through holes and vapor channel walls; **b)** Zoomed-in view of a single 200 μm column. The white areas are the bare silicon at the bottom of the photomask holes, while the brown areas are coated in photoresist.

Before etching the through holes, the oxide on the backside of the wafer must be removed so that it does not act as an etch stop. This was done by “floating” the wafer on top a bath of 5:1 buffered hydrofluoric acid (see Fig. 6.6). Teflon blocks are used to hold the wafer level as acid is slowly poured into the bath. Once the liquid level reaches the backside of the wafer, surface tension “lifts” the wafer off the blocks while the acid etches away the oxide. This technique (as opposed to simply submerging the entire wafer in BHF) prevents “undercut etching” of the exposed sidewalls of the oxide mask.




Fig. 6.6 – Cross-sectional schematic of wafer “floating” technique used for backside oxide etch while protecting the oxide mask. (Note: wafer thickness greatly exaggerated for clarity.)

At this point, the through-wafer columnar holes were partially etched using a Surface Technology Systems (STS) Advanced Silicon Etcher. Deep Reactive Ion Etching (DRIE) systems such as this utilize high-frequency inductively coupled plasma, which alternates between cycles of directional silicon etching using SF_6 and protective sidewall deposition using C_4F_8 . This allows for extremely deep etching with very high aspect ratio. Ideally, I would etch all of the through holes to, say, $600\text{ }\mu\text{m}$ (leaving the final punch-through for a subsequent etch, when an underlying handle wafer would protect the STS chuck from damage). Unfortunately, however, there are a number of well-known phenomena that limit etch uniformity across the wafer, and all of them are a factor in my layout. First, non-uniformity in the plasma itself causes faster etching at the edges of the wafer and slower etching in the center—this is known as the *bull's eye effect*. Second, larger holes will etch faster than smaller holes due to reactant transport limits in deep, narrow holes—this is formally known as *Aspect-Ratio Dependent Etching* (ARDE) or more commonly as *RIE lag* (see Fig. 6.7, left). Finally, densely-packed hole arrays will etch more slowly than isolated holes because reactants are more quickly used up around dense hole patterns—this is known as *microloading* or simply the *loading effect* (Fig. 6.7, right).



Fig. 6.7 – Two causes of etch-rate non-uniformity. Smaller holes will etch slower due to *RIE lag* (left), while differences in hole array density will result in non-uniform etching from *loading* (right).




Fig. 6.8 – Schematic wafer cross-sections for Floor process flow (Steps 9 – 11).

Because I wanted to minimize the amount of etching while using a handle wafer (since it is slower and does not allow me to see the backside of the device wafer), I simply etched without a handle wafer until *the first* hole punched through the wafer. To minimize etching of the chuck once the first hole did punch through, very short etching steps (< 5 min) were done toward the end of the process. As expected, the first punch-through was a large feature at the edge of the wafer (actually part of a lithography test structure). This “warning” punch-through of the test structure concluded the first DRIE process, and the remaining photoresist was removed in an 80°C PRS-3000 bath. Partially etched through holes can be seen in Fig. 6.9, while characteristic device features after the photoresist strip can be seen in Fig. 6.10.




Fig. 6.9 – Partially etched $400\ \mu\text{m}$ holes and masked vapor channel walls (Step 10); **b)** Zoomed-in view of a single $400\ \mu\text{m}$ hole with the objective focused on the bare silicon at the bottom of the hole.




Fig. 6.10 – Micrographs of the characteristic features of two devices after the thick photoresist soft mask has been stripped, revealing the oxide hard mask (Step 11): **a)** $800\ \mu\text{m}$ through holes and vapor channel walls; **b)** $200\ \mu\text{m}$ through holes and vapor channel walls. (Note: the lighter tan areas are bare silicon, while the darker brown areas are the oxide mask.)

Next, the wafer was cleaned in Piranha and then reversibly bonded to an oxide-encapsulated handle wafer. This is done by coating the handle wafer with 2 μm of G-line photoresist (with no soft bake) and then firmly pressing the device wafer onto the handle wafer (with the photoresist acting as “glue”). This wafer stack is then placed in a 90°C oven under vacuum for approximately 30 minutes to evaporate the solvents and cure the bond. During the next DRIE step, STS is again used to etch the channels (thus also defining the exterior walls of the columns), while simultaneously punching through the remaining columnar holes. The handle wafer is required during this punch-through so that continued etching through the already-completed holes does not damage the STS chuck. Because a very large area of the wafer is etched during this step, the bull’s eye effect is more pronounced, and etch uniformity is poor. Devices near the outside edge of the wafer have channel depths of around 250 μm , while devices near the center have channel depths closer to 150 μm . However, a large variety of device configurations was desired, this variation was not viewed as a problem. After completion of this second STS etch process, the wafers were stripped of the remaining oxide hard mask using 5:1 buffered hydrofluoric acid and then diced. A completed floor wafer (prior to oxide removal and dicing) can be seen in Fig. 6.12, while a few diced devices are shown in Fig. 6.13.








Fig. 6.11 – Schematic wafer cross-sections for Floor process flow (Steps 12 – 15).




Width: 900 μm ; Pitch: 1250 μm



Width: 500 μm ; Pitch: 2500 μm



Width: 900 μm ; Pitch: 2500 μm



Width: 200 μm ; Pitch: 1250 μm

Fig. 6.12 – Nearly completed 6" wafer (Step 14) and close-ups of four individual devices. Width refers to the *exterior* columnar dimension (*i.e.* a 900 μm column has an 800 μm through hole).




Fig. 6.13 – Diced floor devices: **a)** Column width: 500 μm , Pitch: 5000 μm ; **b)** 200/1250; **c)** 900/1250.

6.2 – Microevaporator Ceiling Fabrication

Ceiling fabrication was also performed on 6” p-type prime silicon wafers with a nominal thickness of 625 μm . This process flow involved only a single lithography step, using a standard I-line photoresist etch mask. Schematic cross-sections of all processing steps can be seen in Fig. 6.14. Wafers were first cleaned with Piranha and HF and then coated with HMDS for photoresist adhesion. The wafers were spin-coated with 1.1 μm of I-line photoresist that was then soft-baked for 60 seconds at 90°C. Using Mask #3 (see Fig. 6.15) in a Karl Suss MA6 Mask Aligner, the photoresist was exposed to define the channel patterns. A post-exposure bake for 60 seconds at 120°C was followed by standard I-line development. Finally the photoresist mask was UV hard-baked using a Fusion M150PC Photostabilizer System. The wafers were then DRIE etched using STS for approximately 10 minutes, leaving channels between 5-20 μm deep. This large variation in channel depth was due mostly to RIE lag and loading effects on a layout with dramatically different channel widths and densities (2-16 μm). However, this effect was actually beneficial because it led to channels with a more “square” cross section (*i.e.* the 16 μm wide channels were closer to 20 μm deep, while the 2 μm wide channels were closer to 5 μm deep. Finally, after this STS etching, the photoresist was stripped in an 80°C PRS-3000 bath, and the wafers were cleaned and diced. Micrographs of all twenty microtextures are shown in Figs. 6.16, 6.17, and 6.18, while a sample of diced ceiling components is shown in Fig. 6.19.




Fig. 6.14 – Schematic wafer cross-sections for Ceiling process flow.

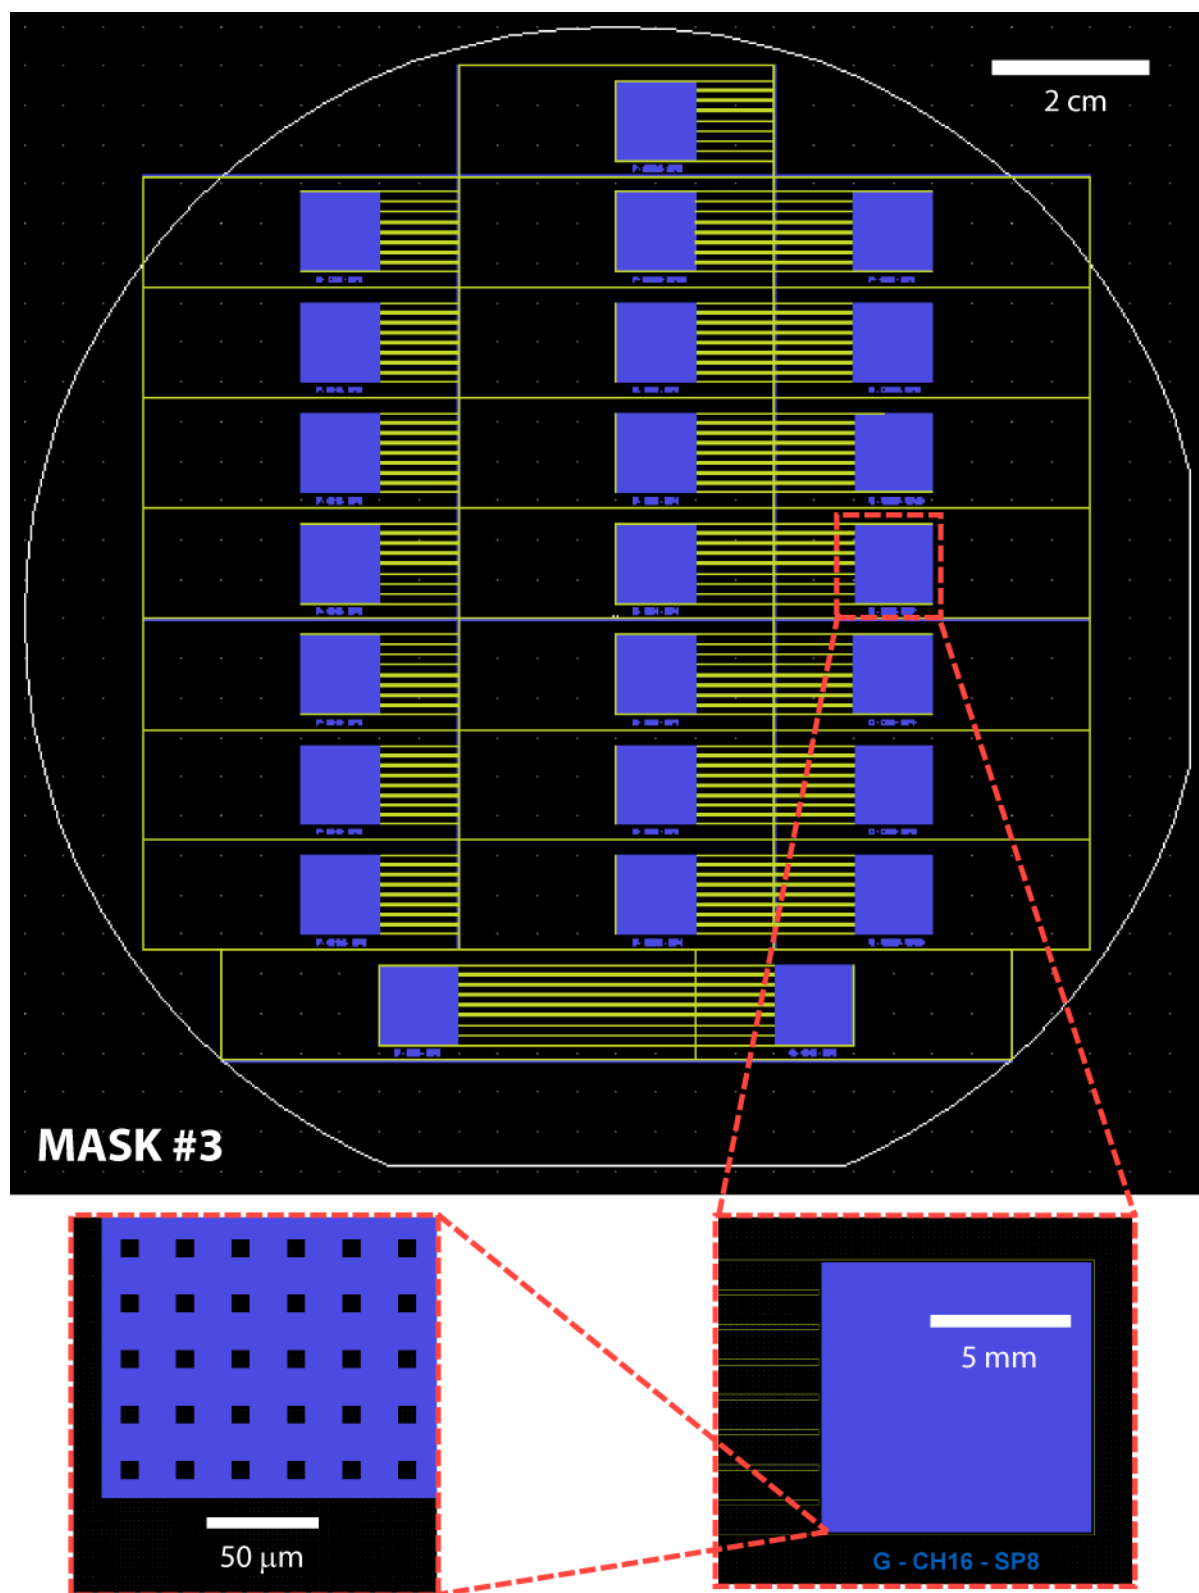





Fig. 6.15 – Photomask for patterning ceiling microtextures. Blue areas are clear on mask and expose/remove positive resist; black areas on mask are chrome and shield resist. (Note: the label “G-CH16-SP8” identifies the zoomed-in device as having a “grid” appearance with 16 μm channels spaced 8 μm apart.)




Channel Width: 2 μm ; Spacing: 2 μm




Channel Width: 2 μm ; Spacing: 4 μm




Channel Width: 4 μm ; Spacing: 4 μm




Channel Width: 8 μm ; Spacing: 4 μm



Channel Width: 8 μm ; Spacing: 8 μm



Channel Width: 16 μm ; Spacing: 4 μm



Channel Width: 16 μm ; Spacing: 8 μm



Channel Width: 16 μm ; Spacing: 16 μm


Fig. 6.16 – Micrographs of all eight parallel channel ceiling microtextures. Magnification and scale are identical in all images.




Channel Width: 2 μm ; Spacing: 2 μm



Channel Width: 2 μm ; Spacing: 4 μm




Channel Width: 4 μm ; Spacing: 4 μm




Channel Width: 8 μm ; Spacing: 4 μm




Channel Width: 8 μm ; Spacing: 8 μm



Channel Width: 16 μm ; Spacing: 4 μm



Channel Width: 16 μm ; Spacing: 8 μm



Channel Width: 16 μm ; Spacing: 16 μm

Fig. 6.17 – Micrographs of all eight orthogonal channel ceiling microtextures. Magnification and scale are identical in all images.




Fig. 6.18 – Micrographs of all four oblique channel ceiling microtextures. Magnification and scale are identical in all images.




Fig. 6.19 – a) Three diced ceiling components with 1 cm x 1 cm micropatterned surfaces; b) Device floor and ceiling showing placement of microtexture region over evaporator region.

6.3 – Fabrication Minutiae

This section briefly addresses some of the technical difficulties, mishaps, calibration steps, and subsequent process modifications that accompanied the floor and ceiling fabrication. As discussed in the introduction of this chapter, these details were intentionally omitted from Sections 6.1 and 6.2 in order to create a clearer, more linear narrative reflective only of the *final* iteration of the process flow. However, a finalized fabrication process rarely emerges *ex nihilo*,

and a great deal of the microfabrication research experience is dedicated to the trial-and-error development and fine-tuning of a successful process.

6.3.1 – Through-Wafer Etch Requirements

The first “nonstandard” processing step during Floor fabrication was the deposition and patterning of the thick multi-coat photoresist mask for the partial through-wafer etch (see Section 6.1, Steps 7 and 8). Historically, anisotropic through-wafer etching of silicon (for micro-nozzles and inkjet printer heads, for example, [see 6.5]) was accomplished using potassium hydroxide (KOH) with a silicon dioxide or silicon nitride etch mask. KOH etch selectivity between silicon and oxide ranges between 150-350:1, while nitride is not attacked at all [6.6]. Such high selectivity allows for through-wafer etching with a reasonable mask thickness (1.8 to 4.2 μm for oxide). However, KOH etching on $\langle 100 \rangle$ silicon produces the characteristic inverted pyramidal trenches with sidewalls approximately 54° to the surface, and because the slow-etching $\langle 111 \rangle$ planes eventually meet at the “point” of the pyramid, the self-terminating etch depth is defined by the width $2b$ of the mask opening (see Fig. 6.20). Consequently, for a 625 μm thick wafer, the smallest through hole that can be produced using KOH is more than 900 μm across, much larger than all but the largest through holes in my microevaporator layout. Moreover, it was highly preferred that the sidewalls of the columnar through holes be vertical, as well, completely ruling out KOH etching for the $\langle 100 \rangle$ wafers used in this process.




Fig. 6.20 – Anisotropic etching of silicon using potassium hydroxide (KOH).

To achieve vertical sidewalls and through holes as small as 100 μm , DRIE processing was used. For the STS tool in DRIE mode, the etch selectivity of silicon to oxide or hard-baked photoresist ranges from approximately 50-100:1. With this lower selectivity, to etch all the way through a 625 μm thick wafer would require far thicker oxide than is usually deposited using LPCVD. Consequently, I decided to try using a thick photoresist mask deposited by spinning on

multiple coats of 2 μm G-line positive photoresist. Soft-baking between coats minimizes the amount of “settling” (*i.e.* multiple layers compressing to less than the sum of their individual thicknesses). Additionally, I found that G-line resist is much less prone to settling than I-line. On the particular day I was working, the individual coats of G-line were approximately 2.25 μm thick, with 5 coats producing a settled mask thickness of approximately 11 μm . All film thickness measurements were performed using a Nanospec/AFT Model 3000.

6.3.2 – Unexpected Post-Exposure Bake Effects

As should be expected when developing a nonstandard process, this extremely thick photoresist mask presented a host of previously unencountered problems. Two major problems were related to the post-exposure bake (PEB) that I usually perform just prior to development. (A standard PEB usually consists of 60 seconds at 120°C, followed by 6 seconds on a cold plate.) First, during microscope inspection after development, I kept noticing thousands of tiny cracks and bubbles in the photoresist (see Fig. 6.21). These cracks were even visible to the naked eye, but I didn’t initially notice at which step in the development they were occurring. However, one time I happened to notice that *just* as the wafer touched the cold plate, the entire wafer seemed to “frost” over. If I leaned my head over the cold plate, I could even *hear* the tiny cracks forming! I presumed that because of the thickness of the resist, thermal gradients across the layer (as it transitioned from 120°C to room temperature) were stressing the resist to fracture. This doesn’t occur with a single coat of resist because such a thin layer can be quickly and more *uniformly* chilled. Removing the post-exposure bake from the recipe eliminated this problem.



Fig. 6.21 – Cracks in thick photoresist mask due to thermal stresses during post-exposure bake and subsequent cold plate chill.

Second, I also observed dark circular “rings” surrounding the square outlines of the exposed through-hole patterns during microscope inspection after development (see Fig. 6.22). Although I do not know for certain the exact cause of these rings, the effect was greatly diminished when I removed the post-exposure bake from the development process to remedy the thermal cracking issue discussed above (see Fig. 6.23). My belief is that the rings are simply a visual artifact of the out-of-plane curvature at the “mouth” of the mask holes and that this curvature may be broadened during PEB due to reflow of the resist.




Fig. 6.22 – Ring-like appearance perhaps due to curvature at “mouth” of mask holes: **a)** Four 100 μm holes; **b)** Zoomed-in view of a single 100 μm hole; **c)** Cross-section A-A showing broad mouth curvature.




Fig. 6.23 – Diminished ring effect perhaps due to less resist reflow without PEB: **a)** Four 200 μm holes; **b)** Zoomed-in view of a single 200 μm hole; **c)** Cross-section B-B showing tighter mouth curvature.

6.3.3 – Thick Photoresist Mask Exposure Timing

A third complication that arose was the timing of the exposure in the KS Aligner. Such a thick mask obviously requires an unusually long exposure time to fully penetrate the depth of the photoresist. “Standard” exposure times (as per Microlab recipes for a single coat of photoresist) range from 3-5 seconds for a well-conditioned lamp, but calculating the proper exposure is not quite as simple as arithmetic extrapolation; some trial-and-error is inevitable. Additionally, multiple passes through the standard development programs were necessary to fully wash away the dissolved resist. I initially tried a 30 second exposure but found that the photoresist was still very underexposed, with the remaining photoresist measuring between 2000-3000 Å thick and still clearly visible under the microscope and to the naked eye (see Fig. 6.24). Exposure was increased to 45 seconds for the second dummy wafer, and this, too, proved to be just slightly underexposed. Although some of the smallest holes still showed a thin photoresist residue, the large holes were entirely clear of photoresist, except in the corners (see Fig. 6.25). Finally, on the third dummy wafer, the exposure was increased to 55 seconds, and the holes were satisfactorily clean (see Fig. 6.26). This recipe was then used for all of the device wafers.




Fig. 6.24 – Severely underexposed thick photoresist mask after development (exposure time = 30 sec):
a) Array of 800 μm holes with unexposed photoresist visible; **b)** Zoomed-in view of a single 800 μm hole.




Fig. 6.25 – Slightly underexposed thick photoresist mask after development (exposure time = 45 sec):
a) Various sized holes with slight resist residue in corners; **b)** Zoomed-in view of a single 100 μm hole.




Fig. 6.26 – Properly exposed thick photoresist mask after development (exposure time = 55 sec):
a) Array of 800 μm holes; **b)** Zoomed-in view of a single 100 μm hole.

6.3.4 – Unintended Backside Etch Stop

In the very first conception of my process flow, I wanted to avoid using the Tystar LPCVD furnaces because I was not yet qualified to operate them. Instead, I planned to deposit TEOS oxide (tetraethyl orthosilicate) for the hard mask using a Plasma-Enhanced Chemical Vapor Deposition (PECVD) tool called Precision 5000 (shortened to P5000 in Berkeley Nanolab nomenclature). Unlike a furnace, which deposits oxide on all exposed surfaces, this tool deposits oxide *only* on the front side of the wafer. Looking back at Fig. 6.4, had I used P5000 to deposit the oxide for the hard mask, the wafers at Step 8 would *not* have had LTO on the back side and, therefore, would not have required the novel BHF “floating” etch to remove the backside oxide (see Fig. 6.6). Recall that in the final process flow this backside oxide must be removed prior to STS etching because the low etch rate of oxide (~ 20 nm/min) would otherwise create an effective etch stop and prevent punch-through.

Unfortunately, as so often happens in research microfabrication, I made what I believed to be a minor change to my process flow without fully recognizing the peripheral consequences during later steps. In this case, I switched from P5000 to Tystar furnace deposition for the oxide hard mask, without considering the changes to the backside of the wafer. Consequently, for the first batch of device wafers, I started the DRIE through-etching *without first removing the backside oxide* (compare Fig. 6.27, Step 10A to Fig. 6.8, Step 10). Although this had little to no effect on the bulk silicon etching, once the holes reached the oxide deposited on the backside of the wafer, they hit an effective etch stop, and normal punch-through was prevented (compare Fig. 6.27, Step 13A to Fig. 6.11, Step 13). Of course at this time I had completely forgotten that there was oxide on the backside of my wafer, since all of my current process flow diagrams were based upon the earlier assumption of front-side-only P5000 oxide deposition!




Fig. 6.27 – Schematic wafer cross-sections for an earlier iteration of Floor process flow (Steps 10 – 14) when backside oxide was NOT removed prior to STS through-etching. Note that the backside oxide acts as an etch stop during Step 13A and prevents punch-through for most holes.

Discovery of this oversight took some time for two reasons. First, depending on the thickness and uniformity of the reversible photoresist bond, it can be very difficult to tell when a hole has punched through. For larger holes, one can sometimes observe the lighter-colored photoresist at the bottom of a hole using a microscope, which looks very different from the pitch-black bottom of a partially-etched hole (see Fig. 6.28a). However, for small holes (with much larger aspect ratios) it is difficult to properly illuminate the bottom surface (Fig. 6.28b).

Consequently, checking for punch-through requires completely de-bonding the handle wafer, which usually requires an overnight soak in PRS-3000. Then, upon discovering the etch still incomplete, one must repeat the several-hour process of cleaning and re-bonding. This is obviously a time-consuming process, so I wasn't able to quickly or accurately check the progress after each etch step.




Fig. 6.28 – a) Even while still bonded to the handle wafer, the lighter-colored photoresist can be seen at the bottom of some large holes after successful punch-through; **b)** The high aspect ratios of smallest holes, however, often hinder bottom surface illumination. (*Note: the black circles surrounding the white outlines are simply the shadows of the relatively “tall” columns.*)

Second, even *with* the now-forgotten backside oxide, *some* of the holes did punch through, especially those nearer the periphery of the wafer, where etch rates are faster and silicon-to-oxide selectivity poorer. Although the etch time for punch-through seemed unusually long, I had seen poor etch rates with STS before, so I simply continued etching. In hindsight, of course, what was taking so long was etching through the oxide etch stop! With continued etching, however, it eventually seemed as though absolutely no progress was being made for holes near the center of the wafer, particularly the smaller ones. In some cases, a gradient of etching progress could even be seen across a single device (see Fig. 6.29). Such steep gradients were likely due to excess fluorocarbon deposition at the bottom of the holes during passivation cycles of the overly-long DRIE process. Commonly called “grass” among researchers, once these dreaded polymer deposits reach a certain thickness, they become essentially impervious to further etching. The grassy deposits also meant that, once I eventually remembered about the backside oxide, I could no longer complete the punch-through simply by removing the oxide with BHF during Step 15. Luckily, most of the peripheral devices were etched completely through, but I was more or less stuck with the grassy, incomplete ones at the center. Worse yet, in my first (and only) foolish attempt to *manually* punch through some of the remaining grass/oxide membranes using very fine Techni-Tool tweezers, I ended up fracturing many of the column sidewalls (see Fig. 6.30).




Fig. 6.29 – Two examples of steep etch gradients across a single device. The very dark irregular patterns are grassy polymer deposits, while the ragged lighter-colored residues around the hole edges are remnants of the mostly-etched photoresist bonding film. *Note: in both images, the outward radial direction (and thus the direction of more complete etching) is to the left.*




Fig. 6.30 – Broken column sidewalls are clearly visible after a misguided attempt to manually punch through some of the grass/oxide membranes with tweezers

Chapter 7: Experiments

The previous chapter described the operational concept of the open loop columnated evaporator and the overall experimental philosophy and method. This chapter describes how that philosophy was put into practice in a laboratory setting to obtain meaningful conductivity metrics and other performance data. The chapter is divided into five sections. In the first, I describe the physical experimental setup, with subsections devoted to each of the major components and control/data collection systems. In the second section, I discuss the general experimental strategy, practical methods to address system limitations, and justification of experimental methods. The third section details the procedures for the various tests performed, while the fourth section provides a detailed discussion of the calculations necessary to convert the raw data into meaningful metrics. Finally, in the fifth section, I offer a few concluding remarks on the methodology used and discuss the effects of some of the assumptions.

7.1 Experimental Setup

All of the experiments in this study were performed in the Berkeley Micromechanical Analysis and Design (BMAD) Laboratory at the University of California, Berkeley. The test stand and supporting systems are shown schematically in Fig. 7.1, while photographs of the laboratory setup can be seen in Figs. 7.2 and 7.3. The primary components include clamps to seal the device floor-ceiling combinations, a ceramic heater to supply input heat flux, clamps to hold the ceramic heater against the top surface of device ceiling, an infrared video camera and laptop-run camera control and image analysis software, a syringe pump to provide precision fluid delivery, thermistors to calibrate the IR camera and provide auxiliary temperature data, a laptop-run Labview control system for the thermistors and syringe pump.




Fig. 7.1 – Conceptual schematic of experimental setup.




Fig. 7.2 – Photograph of experimental setup in the BMAD Laboratory.




Fig. 7.3 – Enlarged views of two-layer device and clamps.

7.1.1 – Clamps and Heating System

To allow various floor and ceiling components to be interchanged for testing, a pair of plastic clamps was machined. Since the clamps could not be made of metal (to minimize parasitic thermal conduction), they have to be relatively thick to provide adequate sealing force without significant bending. Dimensioned schematics and photographs of the machined clamps are shown in Fig. 4. The bottom clamp contains a liquid reservoir aligned with the footprint of the columnated evaporator area to allow liquid to be pushed up simultaneously through all of the hollow columns. A 1 mm diameter inlet port was bored through the bottom clamp to allow for plumbing interconnection with the syringe pump, and a 1 mm deep slot was milled to serve as a seat for the silicone gasket. The only notable feature of the top clamp is a 1 cm x 1 cm keyhole aligned with the footprint of the columnated evaporator, which serves as an overhead viewing window for the IR camera. Clamps were machined from both polyoxymethylene copolymer (Delrin) and polyoxybenzylmethylenglycolanhydride (Bakelite). Delrin clamps were used for the majority of the tests, and the material was chosen for its ease of machining, low thermal conductivity, low water absorption (less than 0.25% over 24 hours), mechanical stiffness, and relatively high melting point (175°C). For some of the higher temperature dryout tests (which melted the Delrin clamps), Bakelite clamps were used.

Thermal input is provided by a small (19 mm x 19 mm x 2.5 mm) ceramic resistive heater. The Ultramic 600 heater, manufactured by Watlow Electric and distributed by Therm-X of California, is fabricated from aluminum nitride and has a maximum power output of 16 W/cm² at 120V, a maximum operating temperature of 600°C, and contains an integrated K-type thermocouple. Power is supplied to the ceramic heater by an Agilent E3612A DC power supply.




Fig. 7.4 – a) Dimensioned schematic of device clamps; **b)** Photograph of clamps fabricated from Delrin, with seated silicone gasket; **c)** Photograph of clamps fabricated from Bakelite.

A second set of clamps were used to press the ceramic heater against the overhanging “tab” of the device ceiling to promote heat conduction into the device. Since these clamps are in direct contact with the heater, they had to withstand much higher temperatures than the device clamps and were thus machined from Bakelite. The only noteworthy feature of these clamps is a shallow slot milled in the top clamp to provide an air gap between the heater and the clamp (see Fig. 7.5). This gap limits the contact area between the heater and the clamp and helps to prevent conduction into the clamp, thus causing more of the dissipated heat to flow into the device. Ideally, there would be a second air gap between the bottom clamp and the device to further limit conduction into the clamps, but with such a gap, I was worried that any misalignment during clamping would lead to unequal pressure on the device and likely fracture the brittle silicon.




Fig. 7.5 – Heater clamp design showing milled air gap to limit conduction into top clamp.

7.1.2 – Infrared Imaging System

The primary system for collecting temperature data was a FLIR model A320 infrared video camera outfitted with a 10.0 cm focal length lens and controlled and analyzed using ThermoVision ExaminIR software. The camera is capable of 480×310 resolution with a 2 cm spot size, translating to $120 \mu\text{m}$ per pixel. For the temperature range encountered in this study, the camera accuracy is rated at $\pm 0.2 \text{ K}$. In addition to providing dynamic IR images, the software allows real-time calculation of statistical data and functional manipulation of user-specified variables. Lines, boxes, and irregular domains can be superimposed over the image using standard “Microsoft Paint” drawing tools, and the software instantly plots temperature profiles, calculates mean temperatures and standard deviations over one- and two-dimensional domains. These values are constantly updated and the temporal data is recorded, with a sampling rate of 5–8 Hz, depending on the number and complexity of the calculations. For example, in the screenshot shown in Fig. 7.6, two different temperature profiles are plotted: one across the entire length of the device (Line 1) and the other across the evaporator section only (Line 2). The screenshot in Fig. 7.7 demonstrates additional features such as real-time statistics for both linear

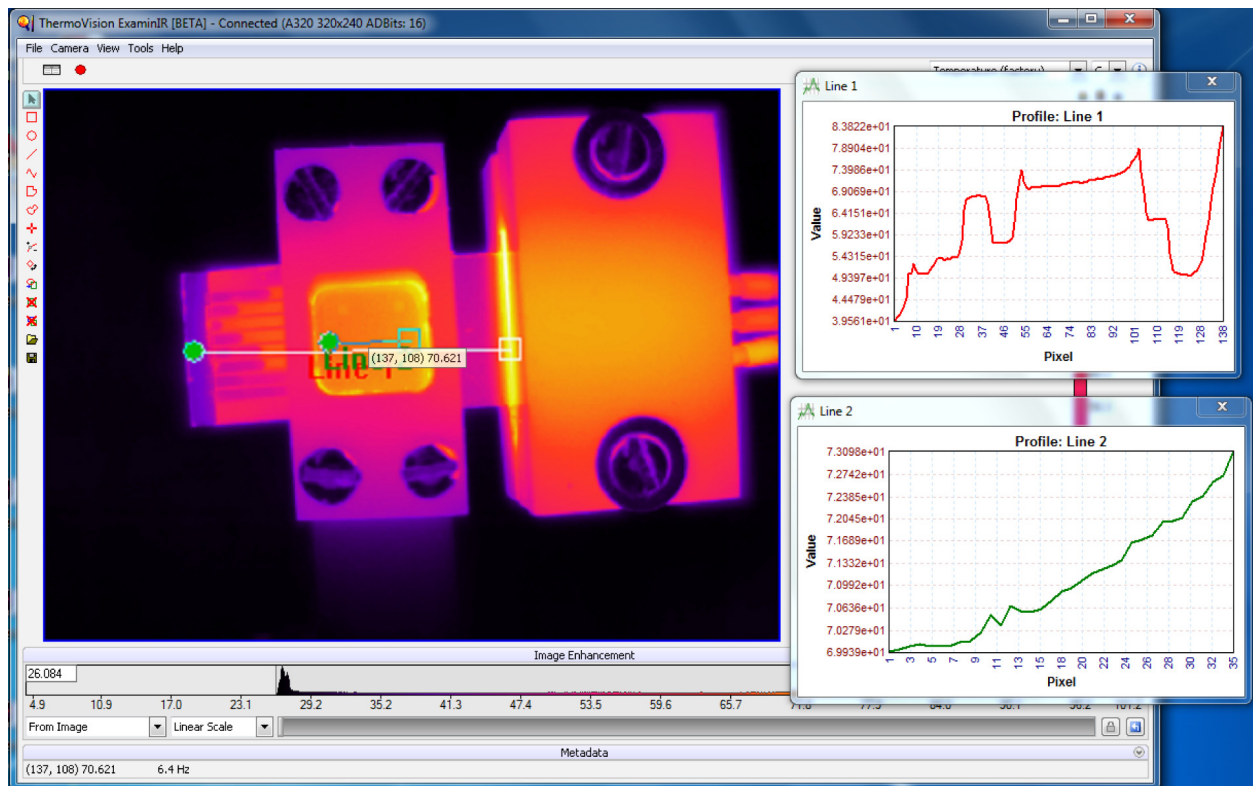


Fig. 7.6 – Screenshot of ExaminIR software plotting dynamic temperature profiles in real time.

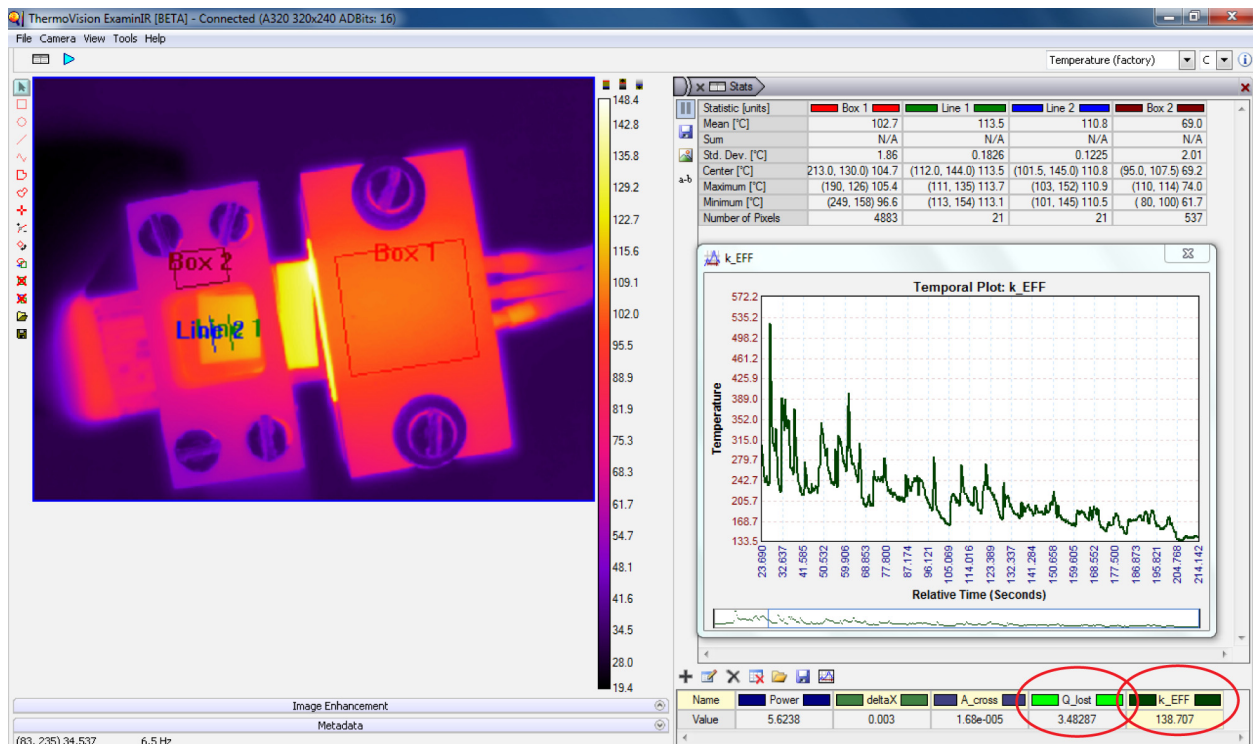


Fig. 7.7 – Screenshot of ExaminIR software calculating statistical data (table in upper right), user-defined variables (circled in red in lower right), and graphing time dependent values (inset plot).

and areal domains, as well as user defined variables that are themselves functions of the statistical values. In this image, for example, the average temperature across two boxes is used to calculate the convective losses from the clamps (“Q_lost”), while the difference between the mean temperatures along two lines drawn on the evaporator surface is used to calculate the temperature gradient, and thus the instantaneous effective thermal conductivity (“k EFF”). Lastly, the inset plot in Fig. 7.7 also demonstrates the ability of the software to record and plot temporal data, such as fluctuations in effective thermal conductivity as the dissipated power is varied in time. All of these tools greatly streamlined the data collection and analysis and allowed much more accurate calculations of the relevant thermal transport. For further discussion of data collection and the precise methods used, see Section 7.3.

7.1.3 – Thermistor System

Although infrared imaging was used for all of the data collection in this study, the camera itself must be calibrated for factors such as surface emissivity and ambient radiation. It is also useful for monitoring surface temperatures not within view of the camera. For this purpose, I implemented the same thermistor-based data collection system used in my Master’s Thesis. The commercial-grade leaded NTC thermistors (model QTMB-14C3, manufactured by Quality Thermistor) have a maximum bead diameter of 0.038" for extremely fast thermal response ($\tau_{\max} = 1$ sec, dissipation constant = 1 mW/°C, both in still air). They have a nominal resistance of 10 k Ω at 25°C, an operating range of -55–125°C, and a tolerance of $\pm 0.2^\circ\text{C}$ from 0-70°C. The thermistor leads are soldered to a 9-pin D-Sub connector, which is then interfaced with ribbon cable via crimp-type connectors to transmit the sensor signals to the data acquisition system

Thermistor data and heater control is accomplished using a high-performance USB-based multifunction DAQ from National Instruments (model 6221, see Fig. 7.8) utilizing a custom virtual instrument constructed using Labview 9 (see next section for further details). The system works as follows: A 5V bias is applied in sequence across each of the thermistors. A series of Wheatstone bridge circuits are used to perform comparative readings to accurately measure the voltage drop across each thermistor (subscript th). This voltage drop is then converted into a corresponding resistance by the following relations:

$$V_{th} = \frac{R_1}{R_1 + R_4} V_0 - \frac{R_{th}}{R_S + R_3} V_0 = \frac{\frac{R_1}{R_4} - \frac{R_{th}}{R_3}}{\left(1 + \frac{R_1}{R_4}\right) \left(1 + \frac{R_{th}}{R_3}\right)} V_0 \quad (7.1)$$

Letting $R_1 = R_3 = R_4 = R$:

$$V_{th} = \frac{1 - \frac{R_{th}}{R}}{2 \left(1 + \frac{R_{th}}{R}\right)} V_0 \quad (7.2)$$

Rearranging and solving for R_{th} :

$$R_{th} = \frac{V_0 - 2 V_{th}}{2 V_{th} + V_0} R \quad (7.3)$$

Sensitivity is maximized by using bridge resistors that correspond to the nominal resistance of the thermistors ($R = 10 \text{ k}\Omega$). Using the standard Z-curve coefficients, this measured resistance is converted to its corresponding temperature using the relation below, which is then recorded on a strip chart.

$$\frac{1}{T} = A + B \ln R + C (\ln R)^2 + D (\ln R)^3 \quad (7.4)$$

$$\begin{aligned} A &= 0.001116401465500 \\ B &= 0.000237982973213 \\ C &= -0.000000372283234 \\ D &= 0.000000099063233 \end{aligned}$$

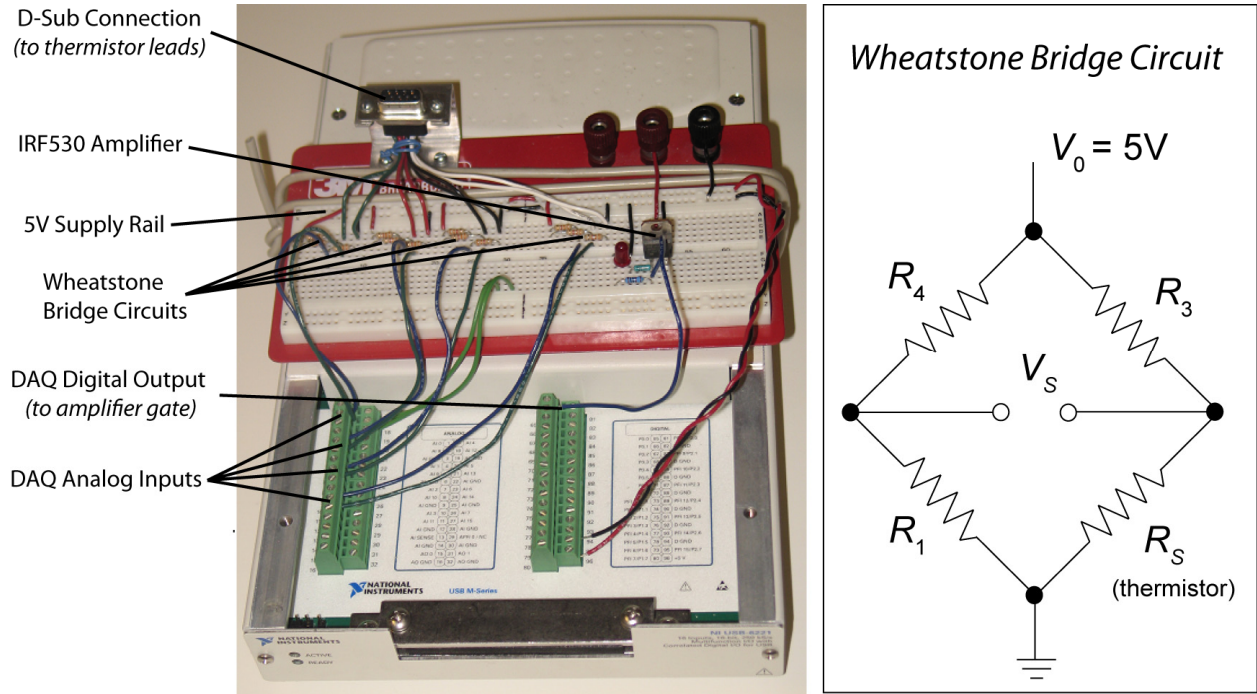


Fig. 7.8 – National Instruments 6221 DAQ with Wheatstone bridge circuitry for converting differential thermistor voltages into corresponding temperature readings.

7.1.4 – Syringe Pump and Control System

The servo-controlled syringe pump used in this study was designed by Thomas H. Cauley, a former BMAD labmember. It can be controlled using either velocity or pressure feedback (using an integrated load cell) and can provide a wider range of flow rates beyond the velocity limits of the gearing by simply changing the diameter of the syringe. In this study, a 10 mL syringe with an internal diameter of 14.2 mm was used. The syringe pump has a USB output

and can also be controlled using Labview; a screenshot of the custom virtual instrument front panel is shown in Fig. 7.9. This virtual instrument combines both the syringe pump control and the thermistor data collection system discussed in the previous section into one program.

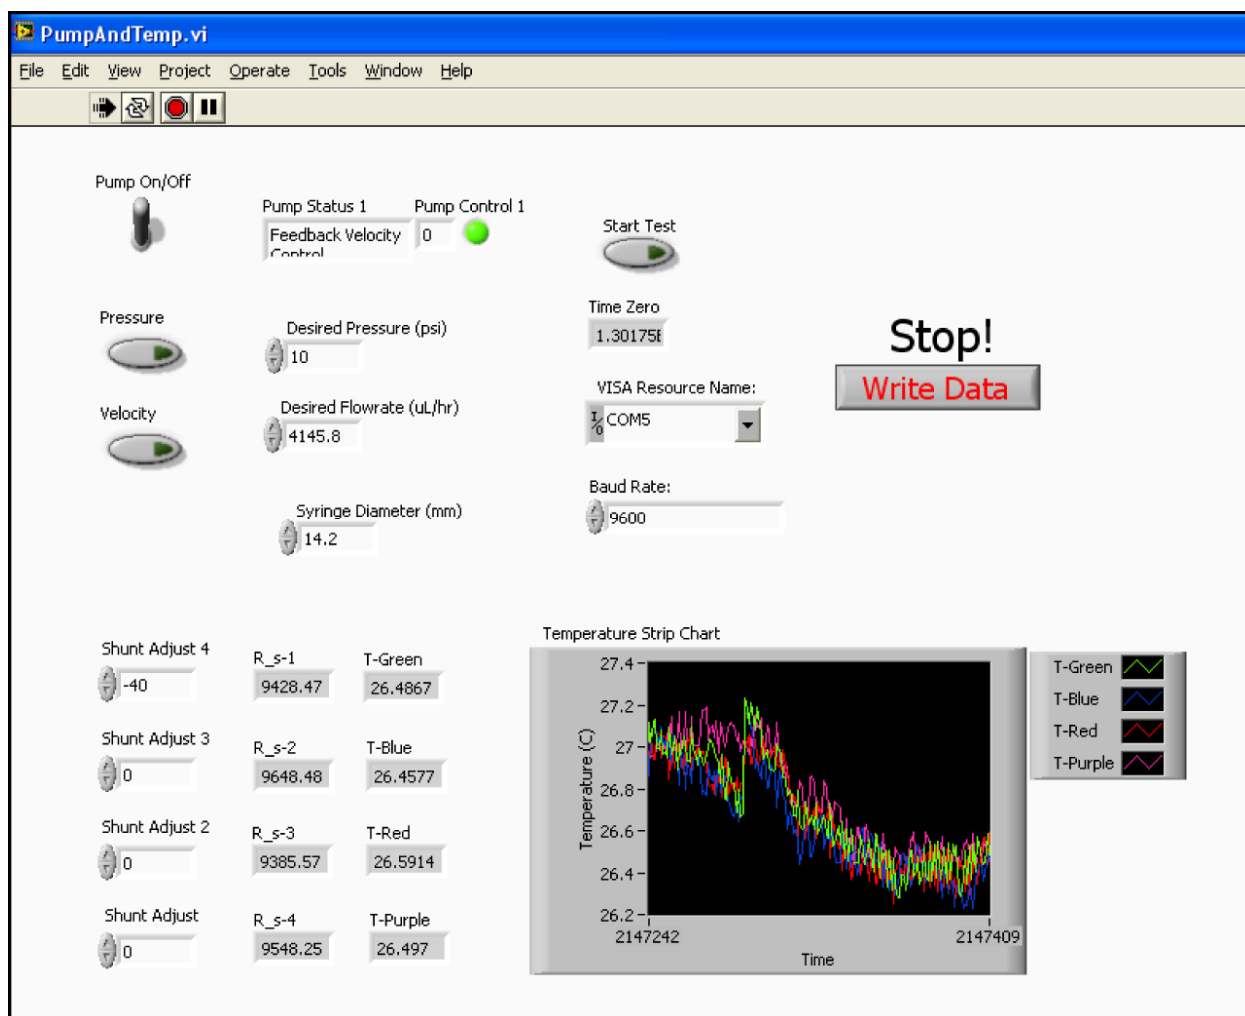


Fig. 7.9 – Screenshot of custom Labview virtual instrument front panel. The single program controls the syringe pump and collects temperature data from four thermistors for camera calibration and monitoring temperatures of surfaces not visible to the camera (e.g. undersides).

The syringe pump is connected to the reservoir fill port on the bottom side of the bottom device clamp using 360 μm steel tubing and commercial MEMS interconnect from Upchurch Scientific. The Upchurch “Nanoport” components (P/Ns C-360-100 and C-360-300) have an internal bore for 360 μm diameter tubing and a 1 mm diameter nipple insert to interface with the through-hole of the bottom clamp. The “female” fitting was permanently epoxied to the bottom clamp using E-6000 industrial adhesive, while the “male” fitting was simply thread-tightened turning “wet” testing and unscrewed and detached during “dry” testing. Even without full epoxy encapsulation, no leakage was ever observed at the Nanoport connections, and such a detachable setup made floor-ceiling interchange much easier.

7.2 Experimental Strategies and Limitations

One of the most difficult challenges in any heat transfer experiment is controlling and estimating the “parasitic” losses, that is, thermal energy dissipated along paths other than that being intentionally measured. In this study, the calculation of the effective in-plane thermal conductivity ultimately relies on two external measurements: the total heat transferred across the device and the local temperature gradient. However, the heat transferred across the device is *not* equal to the power dissipated by the ceramic heater (*i.e.* current times voltage) because not all of the heat from the heater is ultimately transferred axially across the device. Some of the heat is lost through natural convection from the outer surfaces of the various clamps, and some is lost due to natural convection from the exposed silicon surfaces. Ideally, the entire test setup would be placed in a vacuum chamber to eliminate parasitic convection from these surfaces, but if that is not practical (as in this case), such losses can also be minimized by using nonconductive materials and a few design modifications. One downside to using insulating materials, however, is that the system takes longer to reach steady state when the input power is changed. In addition, thermal losses to the heater clamps can be further diminished by decreasing the contact resistance to the *device*, thus increasing the *relative* contact resistance to the clamps. This can be done by adding a thin layer of thermally conductive paste between the heater and the silicon or, conversely, by minimizing the contact area between the clamp and the heater (by machining an air gap, for instance; see Fig. 7.5).

Another potential problem is the *variation* in contact resistance between the clamps and the silicon components when floors and ceilings are interchanged. These variations can be minimized (relative to the total dissipated power) simply by maximizing the amount of heat entering the device (as discussed above), but attempting to create identical test conditions over and over again is quite difficult and/or extremely time consuming. This problem is exacerbated by the fact that, for a system designed to conduct heat efficiently, contact resistances may be a significant contributor to the total resistance. Small variations in the tightness of the screws, the heater-tab overlap (see Fig. 7.10), or the thickness of the thermal paste, for instance, will change the contact resistance, which will in turn change the total flux entering the device. A torque wrench would increase the precision of the clamping force, but only to a point. Moreover, as the system heats up, differences in thermal expansion between the plastic clamps, the steel screws, and the silicon device layers will cause the clamping pressure (and thus the contact resistance) to change throughout the test. In summary, assuming an input flux based solely on the power dissipated by the ceramic heater is a poor technique for calculating thermal conductivity.

An alternative strategy is to attempt to “zero out” the disparities between experiments by analyzing device performance variations *within a single setup*. In other words, as opposed to trying to recreate identical test conditions for every clamped floor/ceiling combination, a baseline test can be run on a dry device (*i.e.* with no fluid flow) to determine the device conductivity for that *particular* setup. This baseline performance can then be used as a reference when analyzing the performance at various heat inputs and liquid flow rates. In theory (*i.e.* if the setup were identical every time), the baseline “dry” performance of all the devices would be more or less equal and very close to the intrinsic solid conductivity of silicon, around 130 W/m·K. This also provides a quick confirmation of the gradient-flux measurement technique for calculating effective in-plane thermal conductivity.

Heater Stack Cross Section (Device Clamps not shown)




Fig. 7.10 – Variations in the overlap area between the device ceiling and the ceramic heater will alter the thermal resistance into the device and thus change the input heat flux.

After “normalizing” for setup variations, the next issue is estimating the parasitic losses. After all, without knowledge of how much heat is lost to parasitic convection, it’s impossible to know how much heat is transferred across the device. In many cases, the best approach is to create an FEM model of the device and clamps, solve for the temperature distribution, and then numerically integrate the flux over the exposed surfaces. However, this would have to be done for each power level and then subtracted out from the power data later. Fortunately, since this study already utilizes real-time IR imaging, it is much simpler (and more accurate) to estimate the parasitic convection using the *actual* surface temperatures recorded in the experiment. In this study, I used the ExaminIR software to calculate average surface temperatures on the heater clamps and device clamps, which I then used to determine appropriate free convection heat transfer coefficients and the corresponding convective heat losses. The convection coefficient will, of course, vary over the surface of the clamps. This is due not only to temperature variation over the surface, but also to surface orientation. Natural convection is a consequence of buoyancy and thus gravity-dependent (*e.g.* convection from the top surface is much stronger than convection from the bottom surface). However, to first order, this method provides a very reasonable estimate of the convective losses and, moreover, can be calculated and subtracted from the total dissipated power in real time. Also, it should be pointed out that, in general, there was very little temperature variation over the surfaces of the clamps (typically the standard deviation was less than 5°C), so the use of constant surface temperature is justified. For a more detailed discussion of parasitic convection estimation, see Section 7.4.

7.3 Experimental Procedures

Having discussed in the previous section the general experimental strategies and limitations, I will now describe in some detail the specific experiments performed in this study. A few of these experiments were designed simply to test and prove the validity of the aforementioned experimental strategies, while others were geared specifically toward the project goal of performance trend identification and evaporator optimization. For some descriptions, I found it easier and more efficient to explain experimental procedures using “cookbook-style” step-by-step instructions, while other times, a simple explanatory paragraph is used.

7.3.1 –Temperature Measurements of IR-Transparent Silicon

Single crystal silicon is largely transparent to infrared radiation; for this study, this was both a blessing and curse. On the one hand, IR transparency allows us to see and measure the temperature of the water and water vapor *inside* the optically opaque silicon devices. We can see the liquid film as it emerges from the interior columns, spreads over the microtextured evaporator surface, and finally enters the vapor channels (see Fig. 7.11). On the other hand, this also means that radiation from outside (*i.e.* beneath) the device can “pollute” the overhead IR images captured by the camera. From a practical standpoint specific to this study, it means that silicon surface temperatures will appear lower than they actually are due to underlying radiation from reservoir (for the evaporator area) and from the table (for the overhanging silicon).

There are two crucial areas for which the silicon surface temperature must be known as accurately as possible. Most importantly, I need to monitor the maximum temperature of the device that is in direct contact with the device clamps to make sure that it does not approach the melting temperature of Delrin. This point is always located on the heater side of the device clamps, just adjacent to the clamps (see Fig. 7.12a). During my first exploratory experiments (before I realized that the underlying “cold” radiation was artificially lowering the observed silicon temperature), I melted the device clamps and had to machine a new set. The second crucial area, of course, is the surface of the actual evaporator region, visible through the viewing window of the top device clamp. This is where data is collected to calculate temperature gradients and, ultimately, the effective thermal conductivity.

Luckily, there is a very simple solution to this transparency dilemma: deposit on the silicon a layer of material that is opaque to infrared. The layer must be thin, so as not to dramatically alter the conductive properties. Had I realized this potential problem during device fabrication, I would have simply evaporated a few hundred nanometers of metal on the backside (*i.e.* upper surface) of the device ceilings. However, since it *is* sometimes useful to see through the device, I would have probably patterned the metal film with windows at certain locations to allow fluid visualization. In any case, having not considered this issue *a priori*, I considered spray painting a series of stripes on the devices to block the infrared radiation. This would likely have worked but would be messy and somewhat irreversible. In addition, many paints do not adhere well to very smooth, metallic surfaces, and the paint may have tended to bead up or degrade at elevated temperatures.

The solution I settled upon was to use small pieces of 40 μm thick Kapton tape in the critical monitoring areas (see Fig. 7.12b). While obviously not as thin as evaporated metal, Kapton tape is completely opaque to IR and compatible with high temperatures. Moreover, it




Fig. 7.11 – Series of six sequential IR images demonstrating how the transparency of silicon allows visualization and temperature measurement of interior fluid. **a)** Superimposed image of the enclosed device showing column locations; **b)** Liquid emerging from column locations; **c)** and **d)** Liquid spreading over entire microtextured evaporator surface; **e)** Fluid entering vapor channels; **f)** Fluid emerging from device.




Fig. 7.12 – a) Location of the maximum temperature encountered by the Delrin clamps; **b)** IR-opaque Kapton tape was attached to the devices at the locations requiring accurate temperature measurement.

does not contribute significantly to the lateral conductivity of the silicon and is easily attached and removed. One issue that remained, however, was whether the conductive resistance across the *thickness* of the Kapton tape would impose a significant temperature drop between the silicon surface and the upper surface of the tape, thus negating the original goal of accurately measuring the silicon surface temperature. To test this effect, I attached different thicknesses (*i.e.* different number of layers) of Kapton tape to the ceramic heater and then compared the surface temperatures of the various Kapton stacks as indicated by the IR camera (see Fig. 7.13).

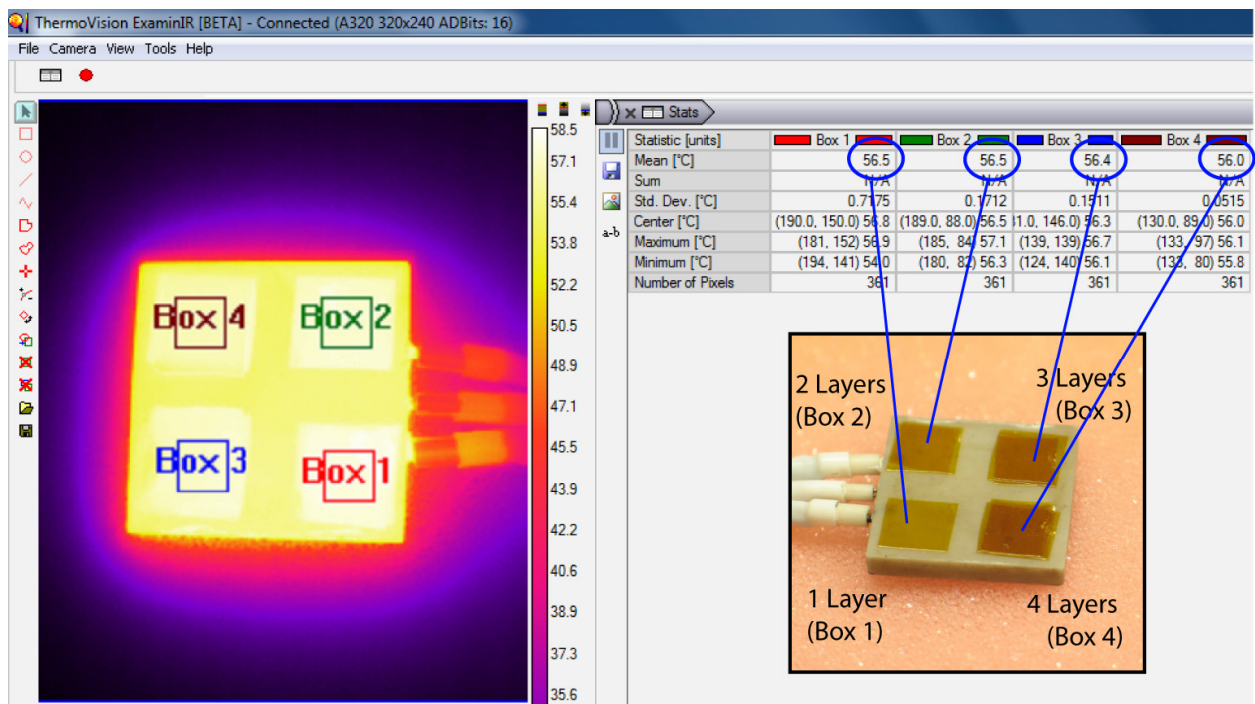


Fig. 7.13 – Limited surface temperature variation over four different thicknesses of Kapton tape indicate that the temperature drop across the thickness of a single piece of tape is negligible.

As highlighted by the blue circles, there is only a 0.5°C temperature difference between the one- and four-layer stacks. Since the four-layer stack has four times the conductive resistance, the temperature drop across the thickness of a single piece of tape is close to or less than the rated precision of the camera ($\pm 0.2\text{K}$). This justifies the use of Kapton tape to accurately measure surface temperatures of the silicon. NOTE: Although the image indicates that the bare (*i.e.* un-taped) surface of the ceramic heater is significantly *cooler* than the taped regions (compare yellow to white on the temperature color bar), this is clearly physically impossible and is, in fact, due to disparities in the surface emissivity of the materials.

7.3.2 – Concept Verification: Solid Conduction vs. Vaporization

Before even attempting to calculate convective losses and obtain meaningful metrics regarding thermal conductivity, I needed to first simply confirm the basic operating principle of the device and test stand, as well as gain experience with the data collection and auxiliary systems. Essentially, what I wanted to verify, first qualitatively then quantitatively, is the concept illustrated in Fig. 5.4. That is, would vaporization of the working fluid and the resulting in-plane convection of the associated latent heat dramatically reduce the temperature gradient across the evaporator? To test this concept, I utilized the ExaminIR software to plot the temperature profile on two linear domains, one that spanned the entire width of the evaporator and one smaller segment centered in the middle of the evaporator (see Fig. 7.14), as the devices were operated under various conditions. In addition, I collected areal statistics on two square domains covering two different percentages of the evaporator. As I incrementally increased the dissipated power and then engaged the syringe pump (once surface superheats were sufficient for vaporization), I measured the temperature gradient across each of the linear domains, as well as the mean, maximum, and minimum temperatures and standard deviation of all of the domains.

Although without calculating the convective losses, it is impossible to know the effective thermal conductivity (see Section 7.4), a change from solid-conduction-dominated heat transfer to vaporization-convection-dominated heat transfer should be evidenced by a significant change in the temperature gradient. In addition, as the evaporator approaches the idealized “thermal ground plane” (*i.e.* an isothermal surface), the areal domains should show a significant decrease in the standard deviation of the temperature distribution. Results are reported in Chapter 8.

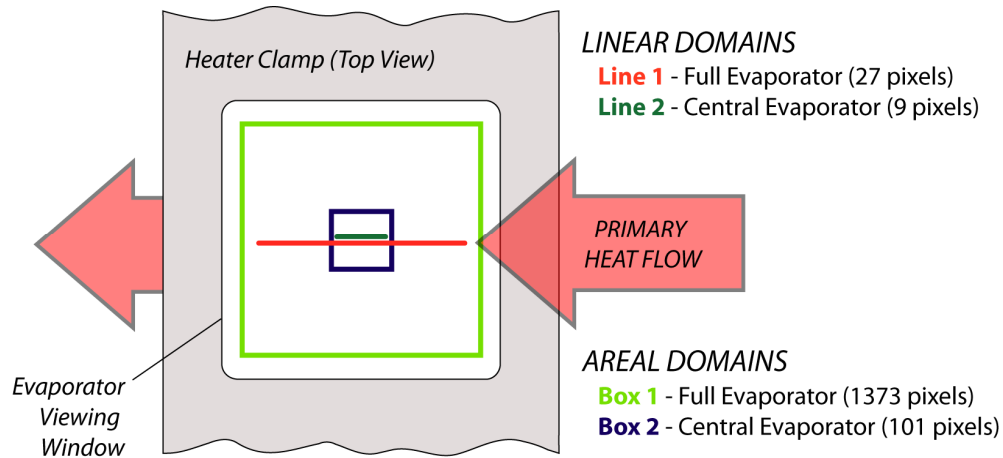


Fig. 7.14 – Statistical domains monitored during initial concept verification tests.

7.3.3 – Primary Conductivity Tests

Having discussed the various preliminary experiments designed to first test the data collection methods, confirm the fundamental concept of the test platform, and determine suitable operating conditions, I will now describe the procedures associated with the primary experiments of this study. Specifically, these experiments were designed to reveal any general performance trends with regard to the columnation scheme and operating conditions and ultimately provide future design guidance to the μ C-LHP program as it transitions from in-plane parallel-channel wick designs to columnated CPS out-of-plane wick designs. Moreover, even if clear performance trends are not confirmed, these experiments will at the very least provide the μ C-LHP team with some valuable baseline data regarding performance limits in future designs.

As discussed in Section 7.2, one of the primary strategies to circumvent variations in input flux and parasitic losses between different tests is to compare a device's performance during active vaporization to its baseline performance without any fluid flow. Thus the first step in every test was to perform a baseline “dry” conductivity test. The steps are as follows:

1. Select a floor and ceiling pair.
2. Align the floor and ceiling components such that the microtextured region of the ceiling sits directly above the columnated evaporator region.
3. Position the two-layer stack in the device clamps with the silicon gasket seated in the milled slot of the bottom clamp and aligned with the columnated evaporator region. Tighten the clamps incrementally, alternating between each of the screws to keep the pressure equally distributed over the device. Clamps were deemed sufficiently tight (and no leakage was observed) when the top clamp was just visibly bowed.
4. Place clamped device in test stand and center the ceramic heater on top of the overhanging ceiling tab. Position heater clamps directly overtop heater and tighten. With no compressible silicone gasket in the heater stack, it was generally pretty clear “by feel” when the clamps were sufficiently tight.
5. Attach Kapton tape to the crucial measurement areas (see Section 7.3.1)
6. Activate the autofocus feature of the camera.
7. Adjust the *absolute* position/angle of the superimposed measurement lines and boxes on the software if necessary, but do not change the *relative* position of the measurement lines to one another (*i.e.* select and move the lines as a group, so as not to alter their spacing, since we want to keep the same Δx for all experiments). Make sure that the measurement lines on the evaporator are contained entirely within the region covered by Kapton tape.
8. Begin collecting IR image data and initiate real time plotting of effective thermal conductivity as a function of time.

9. Set the ceramic heater voltage to 14.0V on the DC power supply and wait for the system to reach steady state. Starting from room temperature, this takes around 30 minutes. NOTE: This particular power level was chosen because it was close to the maximum level that still allowed the devices to be handled comfortably with bare hands. This made subsequent attachment of the fluidic interconnect easier.
10. Once the device has reached steady state, pause camera and save data. Three sets of data files were collected: 1) Portable Network Graphics (.PNG) screenshot of entire computer display; 2) Comma Separated Values (.CSV) file of IR image (in case more detailed analysis of non-plotted data is desired later); and 3) Text File (.TXT) of the statistical data (*e.g.* mean and standard deviation of superimposed box and line temperatures) and current values of the user-defined variables (*e.g.* total dissipated power, estimated parasitic convective heat transfer coefficient, total parasitic heat loss, effective thermal conductivity, *etc.*).
11. Enter data summary in Excel spreadsheet.

With the dry baseline data established for the given clamping setup, we now look at the device performance as working fluid is introduced and the heat flux is increased.

12. Attach microfluidic interconnect to device clamp reservoir port and turn on syringe pump using Labview virtual instrument. Flow rate for all tests was 5400 $\mu\text{L/hr}$.
13. Increase DC voltage to 19.0V and again initiate real time plotting of effective thermal conductivity as a function of time. This particular power level was selected because it provides sufficient flux to heat the evaporator surface just above the saturation temperature. As discussed thoroughly in Appendix A, phase change is thermodynamically impossible below this level, and in general, 5-10 degrees of superheat are required to initiate vaporization. In the first few tests, I started the ramp-up at lower power to have a more complete plot of k vs. Q , (*i.e.* 15V, then 15.5V, then 16V, *etc.*), but this proved to be extremely time consuming, and I never saw evidence of any vaporization below 19.0V.
14. Again, wait for the device to reach steady state. The increase from 14.0-19.0V also generally took about 30 minutes to stabilize. Ensure that only liquid (no vapor) emerges from vapor channels.
15. Pause camera and again save the three data files and enter data summary in Excel.
16. Above 19.0V, continue the power ramp-up in increments of 0.1V (the minimum step size available for the power supply), and record data at each level upon reaching steady state. When the transition to vaporization begins, “steady state” is, of course, a relative term, but typically a series of easily recognized conductivity “regimes” can be observed (see Section 9.4). In addition, the degree of “steady-statedness” can be quantified by comparing a time-integrated average value of any chosen variable (*e.g.* surface temperature, effective thermal conductivity) to the extrema.

17. Continue to increase the power and record data until the effective conductivity starts to decrease rapidly. This generally begins to occur when the evaporator surface exceeds 115°C, likely due to transition to dryout on a significant area of the evaporator. As the evaporator temperature peaks, be sure to monitor the maximum silicon temperature adjacent to the device clamps (see Fig. 7.12a) and abort test if temperature exceeds 130°C.
18. Decrease DC voltage to 0V and wait a minute or two for the device to cool slightly.
19. Loosen heater clamps and remove ceramic heater.
20. Turn off syringe pump in Labview and detach microfluidic interconnect.
21. Loosen device clamps and carefully remove device components.
22. Evaporate residual water in reservoir and device channels with air gun.
23. Repeat entire process with a new pair of floor and ceiling components.

Even with “easily” interchangeable components, the entire process of testing a single floor-ceiling combination over its vaporization range (for a given flow rate) takes approximately one hour. This is due primarily to the (intentionally) poor conductivity and (unintentionally) large thermal mass of the clamps, which causes steady state to be reached much more slowly than if only the silicon were heated. In fact, one had to be careful not to assume that the highest temperatures observed during power ramp-up were most representative of steady-state conditions. Somewhat counter-intuitively, the progression to steady-state conditions was often non-monotonic for all surfaces. For example, upon raising the voltage from 0-14V in Step 9, the exposed silicon surface often reached temperatures *above* the steady-state value because the resistance through the silicon was so much less than through the Bakelite. This essentially caused heat to initially “bypass” the clamps and “build up” in the silicon until the silicon was sufficiently heated to force heat into the clamps. Naturally, this effect was *much* less pronounced (if present at all) when the incremental voltage increases were small.

7.4 Heat Transfer Analysis

Having discussed the actual experimental procedure, I will now discuss the heat transfer analysis and corresponding data conversion. A complete sample calculation can be found in Appendix D. As introduced in Section 2.1.1, solid conduction is governed by Fourier’s Law. The total heat transferred is directly proportional to the temperature gradient, directly proportional to the cross-sectional area, and inversely proportional to the length. The thermal conductivity k is simply the proportionality constant. Thus, for a finite temperature difference across a finite distance:

$$Q = -kA_{cross} \frac{\Delta T}{\Delta x} \quad (7.5)$$

Rearranging for the purpose of this study and concerning ourselves only with magnitudes:

$$k = \frac{Q}{A_{cross}} \frac{\Delta x}{\Delta T} \quad (7.6)$$

Although devices with different columnation schemes technically have different cross-sectional areas, for this study I am interested in calculating the *effective* thermal conductivity of the evaporator region (*i.e.* as if it were made purely of solid material). Consequently, I have assumed a constant cross section for all devices, namely the entire external dimensions of the cross section (14 mm in width times the 1.2 mm thickness of the double wafer stack). Furthermore, for all tests, I calculated the temperature difference on the evaporator across a distance of 3 mm centered on the evaporator midline. Thus in Eq. 7.6,

$$\Delta x = 0.003 \text{ m} \quad \text{and} \quad A_{cross} = 0.014 \text{ m} \times 0.0012 \text{ m} = 1.68 \times 10^{-5} \text{ m}^2$$

Given the fact that ΔT is measured directly by the IR camera, this leaves Q as the only missing piece in calculating the effective thermal conductivity. Unfortunately, as discussed in Section 7.2, Q is not simply the total power of the ceramic heater because not all of the heat dissipated by the heater actually transfers *across* the device. Much is lost to natural convection from the elevated surface temperatures of the various clamps and exposed silicon surfaces. In other words, approximately:

$$Q_{conducted} = Q_{heater} - Q_{lost} \quad (7.7)$$

Since the heat dissipated by the heater is easily calculated (voltage \times current), the major work involved in calculating the effective thermal conductivity revolves around estimating the heat lost to natural convection. Like convection from any exposed surface, natural (or free) convection is governed by Eq. 7.8 below.

$$Q = h A_{surface} (T_{surface} - T_{fluid}) \quad (7.8)$$

The primary work, of course, is involved in calculating h , and such calculations typically come from empirical correlations based on dimensional analysis. Here, I will not take the time to trace the development of these correlations, but only cite their sources and explain how they were used specifically in this study.

The task at hand, then, is to find the heat lost through natural convection from the exposed surfaces of the clamps. As is apparent from Eq. 7.8, this is first and foremost dependent on the temperature of the exposed surfaces. It is clear from the sample IR image in Fig. 7.15 that the surface temperature varies over the clamps, typically 5-10 degrees from maximum to minimum. However, for a number of reasons (see comments below) I initially chose to calculate convective losses assuming a constant surface temperature for each clamp. To obtain these average temperatures, I used the real-time areal statistics provided by the software for two boxes superimposed over the surfaces of the clamps.

Using these average surface temperatures, along with clamp dimensions and the thermophysical properties of air, I created a local variable in the software to calculate the Rayleigh number, which governs natural convection in much the same way that the Reynolds number governs forced convection. Specifically:

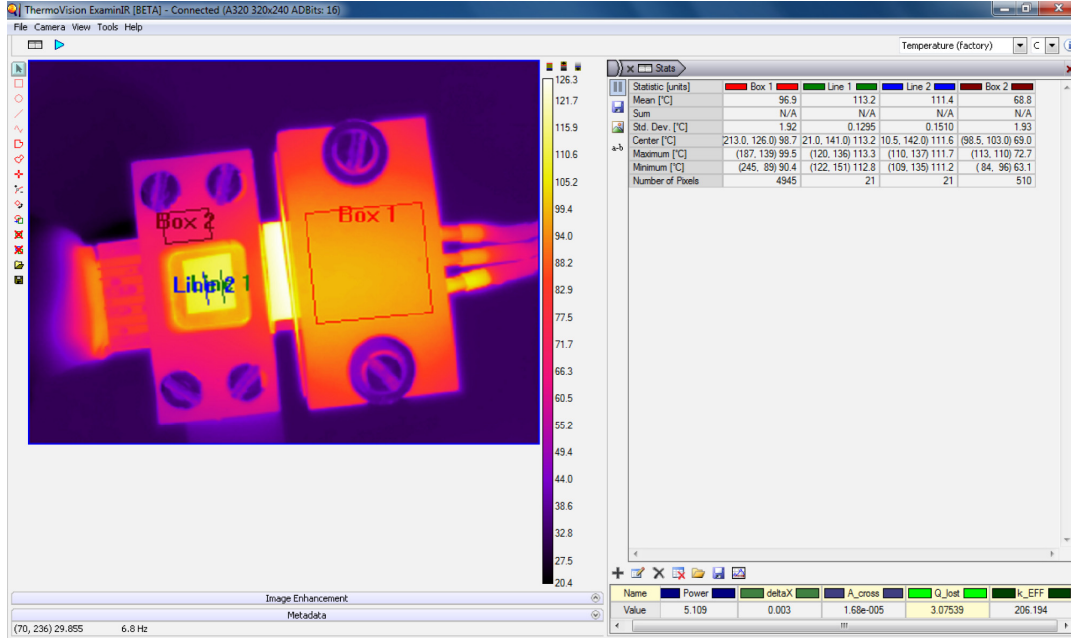


Fig. 7.15 – ExaminIR software screenshot illustrating the superimposed areal and linear domains used to estimate convective losses and calculate effective thermal conductivity.

$$Ra_L = \frac{g \beta (T_{surface} - T_{fluid}) L^3}{\nu \alpha} \quad (7.9)$$

where β , ν , and α are the volumetric thermal expansion coefficient, kinematic viscosity, and thermal diffusivity, respectively, of the surrounding fluid (air), and L is the characteristic length of the surface. The precise definition of L depends on the correlation used. In addition, for ideal gases (such as air at room temperature), the volumetric thermal expansion coefficient β can be simplified to the reciprocal of the absolute temperature of the fluid (see [1]).

$$\beta = \frac{1}{T} \quad (7.10)$$

Although all of these fluid properties are temperature-dependent, the relatively small temperature range encountered in this study, as well as the modest temperature *difference* between the surface and the surrounding air, justifies the standard use of a *film temperature* T_f for calculating thermophysical properties, where

$$T_f = \frac{1}{2} (T_{surface} + T_{ambient}) \quad (7.11)$$

Once the characteristic Rayleigh numbers have been calculated, the average Nusselt number for each of the exposed clamp surfaces is obtained using the empirical correlations in Eqs. 7.12 – 7.14 below. Note that the correlations for horizontal surfaces (Eqs. 7.12 – 7.13) utilize a characteristic length defined as the ratio of the surface area to the perimeter. Also, in

Eq. 7.14, the thermophysical Prandtl number Pr is simply the ratio of the kinematic viscosity of the fluid to its thermal diffusivity, that is, ν/α .

An empirical correlation for top surfaces is (from [2]):

$$\overline{Nu}_L = 0.54 Ra_L^{1/4} \quad L \equiv A/P \quad (2 \times 10^4 \leq Ra_L \leq 10^7) \quad (7.12)$$

An empirical correlation for bottom surfaces is (from [3]):

$$\overline{Nu}_L = 0.27 Ra_L^{1/4} \quad L \equiv A/P \quad (10^5 \leq Ra_L \leq 10^{10}) \quad (7.13)$$

An empirical correlation for lateral surfaces is (from [4]):

$$\overline{Nu}_L = 0.68 + \frac{0.67 Ra_L^{1/4}}{[1 + (0.492/Pr)^{9/16}]^{4/9}} \quad (Ra_L \leq 10^9) \quad (7.14)$$

For the geometry of the clamps used in this study, the area-perimeter ratio definition of characteristic length results in lengths that are about half that of the conventional definition. Furthermore, since the Rayleigh number is proportional to the cube of the characteristic length, the Rayleigh numbers for the horizontal surfaces are one to two orders of magnitude less than those for the vertical surfaces and, more importantly, far below the suggested limits of the correlations. An extensive search of the literature did not yield any appropriate correlations for Rayleigh numbers below 20,000. Although I still utilized these correlations, I also performed a second calculation for comparison that approximates the clamps as spheres. This calculation was performed as a safety check because the Nusselt correlation for spheres (Eq. 7.15) does not have a lower limit for applicable Rayleigh numbers.

An empirical correlation for the average Nusselt number over the *entire* surface of a sphere, where the Rayleigh number uses the diameter as the characteristic length, is (from [5]):

$$\overline{Nu}_D = 2 + \frac{0.589 Ra_D^{1/4}}{[1 + (0.469/Pr)^{9/16}]^{4/9}} \quad (Ra_L \leq 10^{11}) \quad (7.15)$$

Once the average Nusselt number for each exposed surface has been calculated, the average heat transfer coefficient for each surface (h in Eq. 7.8) is obtained using the defined relationship between the Nusselt number and the convection heat transfer coefficient:

$$\overline{Nu}_L = \frac{h_{AVG} L}{k_{fluid}} \quad (7.16)$$

As with the other thermophysical properties, the fluid thermal conductivity k_{fluid} is evaluated at the film temperature of the surrounding air, as defined by Eq. 7.11. Finally, according to Eq. 7.8, each surface area is multiplied by its corresponding h_{AVG} value and temperature difference $T_{surface} - T_{ambient}$ to obtain the convective heat loss from each surface. These losses are then summed to obtain the total convective loss in Eq. 7.7. The entire process of converting the raw temperature data obtained from the IR camera into meaningful conductivity data is summarized in a flowchart in Fig. 7.16 below.

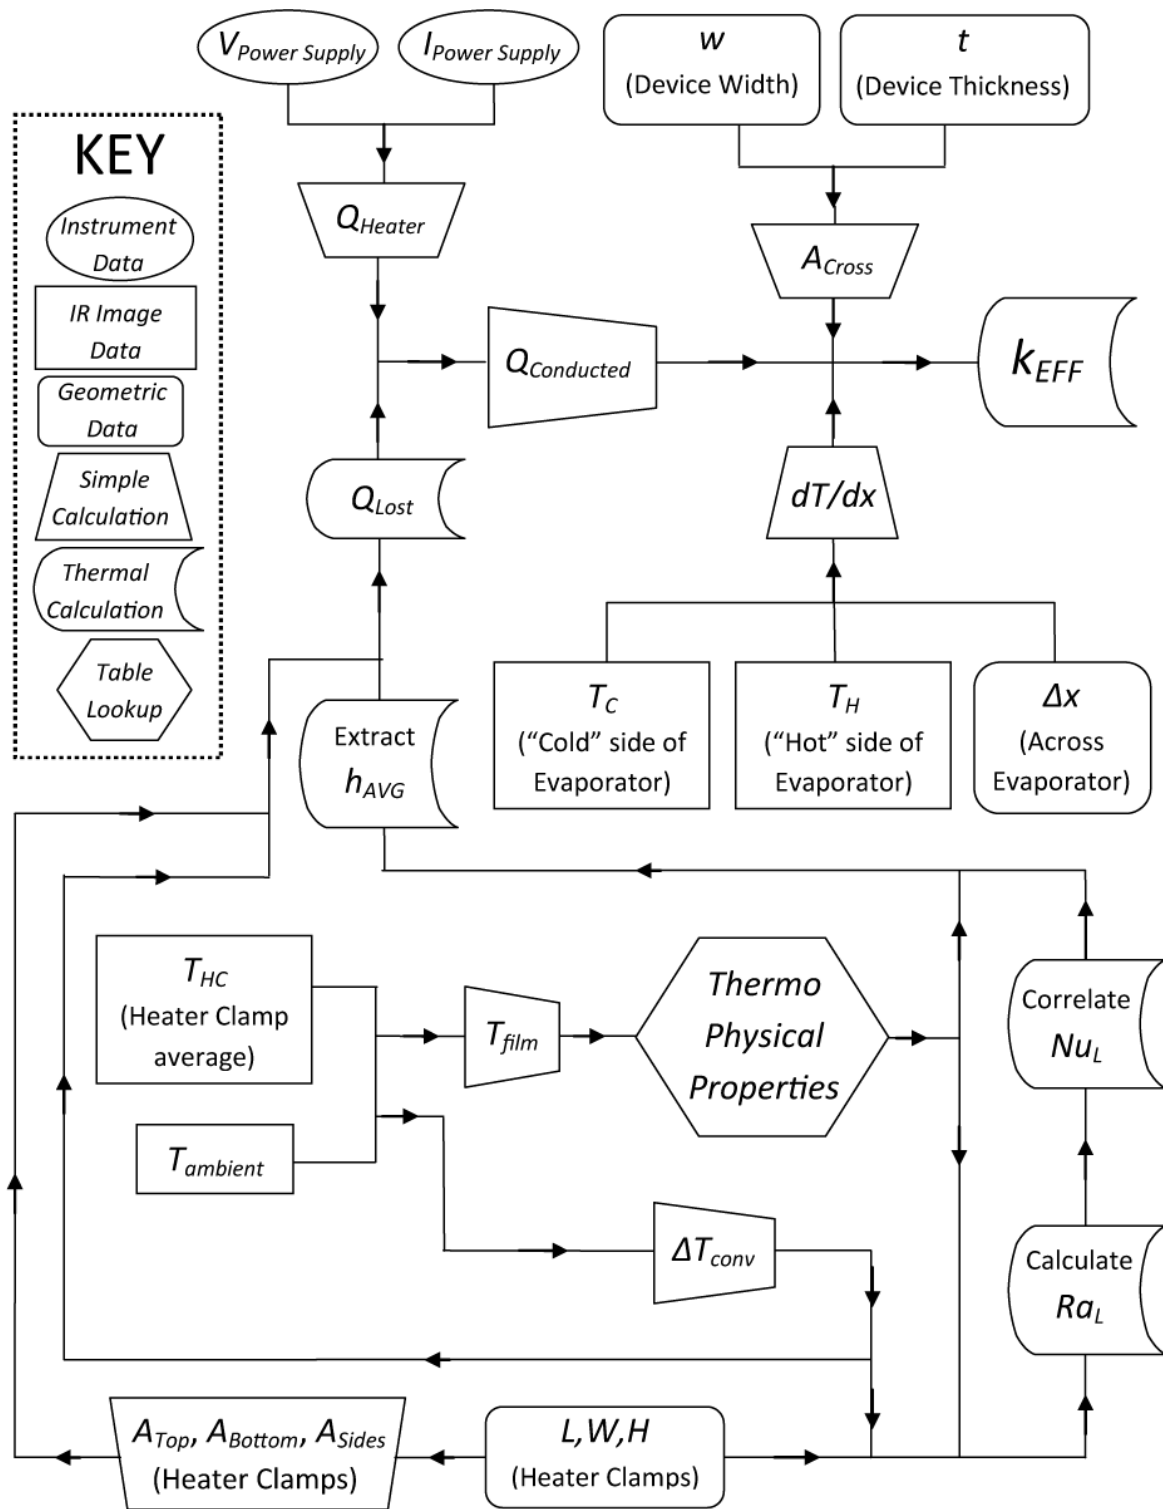


Fig. 7.16 – Flowchart illustrating the calculation process from raw data to meaningful conductivity metrics. (Note: for simplicity, heat losses are depicted for only one set of clamps.)

7.5 Notes on Methodology

At this point, I should comment on a few of the assumptions and shortcomings of this method. The first assumption I made was to assume a constant surface temperature for each of the clamps. As discussed earlier, this is a reasonable assumption because based on statistical data supplied by the ExaminIR software, the standard deviation in temperature over the surface of the clamps was typically less than 5°C. Although clearly an approximate method, if anything, these “characteristic” boxes overestimate the average surface temperatures—since Box 1 is located directly above the footprint of the ceramic heater and Box 2 is located in a region of the clamp that partially overlaps the device. Accordingly, this approximation will tend to *overestimate* convective losses and thus *underestimate* effective thermal conductivity. Additionally, I was careful to make sure that the boxes did not enclose any metal surfaces (screw heads or washers) because their vastly lower emissivity causes them to appear much colder than they actually are (since the camera was calibrated more closely to the emissivity of the clamps and Kapton tape). Ideally, I would have utilized multiple cameras to obtain side and bottom views, as well as more data boxes to better capture variations in surface temperature. Unfortunately, the software only allows real-time data collection over two areal domains, so I was forced to rely on two “characteristic” (but conservative!) boxes. Moreover, at \$10-15K per camera, multiple cameras were simply not financially practical for this study. In order to improve upon this assumption, I made some corrective surface temperature calculations using multiple camera views and empirical correlations (see Section 9.2), but there is bound to be error unless calculations are carried out integrally over the entire surface, an endeavor well beyond the scope of this study.

Secondly, I have not explicitly calculated and subtracted the convective losses from the exposed silicon surfaces, only the clamps. There are two reasons for this, one practical and one technical. Practically speaking, I was not able to collect real-time IR data on these regions for the same reason discussed above: only two areal regions are permitted with the software. While I could have recorded the temperatures *manually* (or use thermistors), this would have greatly slowed an already time-consuming data collection procedure. There was also scientific rationale behind the decision. Obviously, the exposed silicon surfaces are the hottest and, consequently, have the highest convection heat transfer coefficients. However, total convective heat loss is a function not only of heat transfer coefficient and temperature difference, but also of *surface area*. In these experiments, the total surface area of the exposed silicon ($\sim 1.8 \times 10^{-4} \text{ m}^2$) is less than 4% of the total surface area of the clamps. Thus even with much stronger natural convection from these surfaces, their contribution to the total convective loss will be marginal. Additionally, I always included the area of the evaporator window as though it were part of the device clamp (*i.e.* the top surface of the device clamp is treated as a solid rectangle), so I have really only slightly decreased the natural convection from that region, not omitted it entirely.

As a concrete example, in the sample calculation in Appendix D, the total convective heat loss from the *clamps* was estimated to be 2.78 W when the heater clamps and device clamps were at 95.8°C and 64.3°C, respectively, and the evaporator surface was approximately 110°C. In this case, the additional loss from the evaporator surface would be only 0.149 W (less than 5% of the total). Even if the exposed surface of the evaporator region were 150°C (hot enough to melt the clamps!), the total heat lost from *that* surface would be only 0.241 W (still less than 10% of the total). Similar arguments can be made justifying the omission of the convective

losses from the screw and nut surfaces, as well as the exposed lateral surfaces of the ceramic heater itself. Although the heater surface is by far the hottest surface in the system, the exposed area is only $1.9 \times 10^{-4} \text{ m}^2$. Moreover, much of the convective losses from the heater will be “reabsorbed” by the system as the air warmed by the exposed heater flows over the lateral surfaces of the heater clamps. This will artificially raise the film temperature on those surfaces and reduce convection, ideally negating the neglected losses.

The final issue I wish to discuss is the sensitivity of the calculations to intrinsic error in the IR-image-based temperature data. Compared to thermistor-based data (such as that used in my Master’s Thesis, see Appendix B), the IR camera and image analysis software greatly increased the ease and scope of the data collection. Rather than taking temperature measurements at a small number of discrete locations, the IR images provided much more information and consequently presented a much fuller picture of the thermal transport. Even more importantly, thermal imaging offers *much* greater uniformity and control from one test to the next. When I imagine trying to mount, detach, and remount thermistors on each and every device setup (uniformly!), I literally get nauseous. There is also the issue of spatial resolution. Though the thermistor beads are certainly small by most standards ($d = 0.038''$), the $120\mu\text{m}/\text{pixel}$ resolution of the IR camera provides almost an order of magnitude finer resolution. Finally, although the best thermistors are similar in accuracy to the IR camera ($\pm 0.2\text{K}$), that assumes that the thermistor is immersed in an environment of uniform temperature. In other words, yes, a thermistor in a furnace will measure the air temperature to $\pm 0.2\text{K}$, but that level of accuracy must drop significantly when the thermistor is used to measure a *surface* temperature when also in contact with air at a significantly different temperature. Infrared imaging, on the other hand, is designed specifically to detect *surface* radiation and thus measure surface temperature. That is precisely the goal in this study.

Despite all these advantages over thermistors, the certified error of the FLIR camera still presents a bit of a concern. The reason is simple; the most important calculation, k_{EFF} , is *directly* proportional to the measured ΔT across the evaporator (see Eq. 7.6). There are no intermediate calculations to temper the error. Compounding this problem is the small ΔT that characterizes successful operation of the device. Remember, the whole point of utilizing phase change is to try to minimize the temperature drop across the evaporator. An error of 0.2K would not be a big deal if the total ΔT were 10K . However, for most of the experiments, the measured ΔT was between $0.5\text{--}3\text{K}$. Thus 0.2K represents a significant fraction of the total temperature drop.

Contrast this with the general *insensitivity* to error of the convective heat loss calculations. There, a small error in temperature measurement has a much smaller effect on the total heat loss because a change in temperature is first reflected by a small change in the Rayleigh number, which causes a very small change in the Nusselt number, which slightly changes the convection coefficient h , which only then changes the total convected heat. To use the example from Appendix D, if the temperature measured by Box 1 is increased by 0.2K (from 61.4°C to 62.6°C), the increase in the total heat loss from the heater clamps is only 0.0048 W , a change of less than 0.7% . On the other hand, if the evaporator ΔT is increased by 0.2K (from 1.2K to 1.4K), the calculated thermal conductivity decreases from $390 \text{ W/m}\cdot\text{K}$ to $334 \text{ W/m}\cdot\text{K}$, a change of more than 14% .

All is not hopeless, however. There are at least two mitigating factors that help to reduce the error sensitivity in the final conductivity calculations. First, the ΔT measurement is not simply based on temperature measurements at two pixels. I specifically avoided this potential problem by collecting temperature data over finite linear domains (Lines 1 and 2) and then using their *average* temperatures to calculate ΔT . This has the intended effect of randomizing the individual cell errors and thus decreasing the chances of a systemic temperature error in one direction. Second, the short distance between Lines 1 and 2 on the evaporator mean that the detecting cells on the camera sensor array are also quite close together. This tends to decrease the odds of a worst-case scenario (*i.e.* the “hot” pixels are off by +0.2K while the “cold” pixels are off by -0.2K). In other words, even if the measured Line 1 temperature is 0.2K too high, it is likely that the Line 2 temperature is also 0.2K too high, thus minimizing the error in ΔT .

In closing, it must be recognized that heat transfer is one of the hardest quantities to measure accurately (compared to, say, frequency, acceleration, or electric current). Even precise temperature measurements can be extremely challenging in real-world laboratory setups. This makes empirical heat transfer experiments more complicated than most and forces the researcher to find clever design solutions to “zero out” biases and minimize error. Nevertheless, I believe that the strategies and methods used in this study, both from experimental design and procedural standpoints, *do* allow me to collect meaningful data and that the results presented in the next chapter are, subject to reasonable error of course, a legitimate reflection of what is actually occurring inside the microevaporator devices.

Chapter 8: Results and Discussion

This final chapter presents the results of the various experiments described in Chapter 7. It is divided into seven sections. The first section presents graphical data that confirms without a doubt the basic operating principle of the open-loop columnated evaporator. The second section quantifies that performance and describes a clever approach to better estimate the convective losses by improving upon the isothermal clamp analysis detailed in Section 7.4. The third section presents very convincing results of a solid conduction test designed specifically to assess the validity of the heat transfer analysis used throughout this study. The fourth and fifth sections describe the general heat transfer phenomena and performance regimes associated with the various operating conditions of the device, while the sixth section shifts the focus toward maximum performance and design optimization. Finally, in the seventh section, I offer some concluding remarks and summarize the major results. Throughout the chapter, devices are identified according to their floor and ceiling components (*e.g.* B3B4), rather than specifying their columnation scheme and microtexture. See Appendix E for enumeration.

As I did in Chapter 4, I will first present a brief summary of the significant results as a guide to the reader. This is intended to keep the broad goals and major conclusions in mind as the specific results and analysis are presented:

- 1) In the open-loop test setup, the columnated evaporator successfully demonstrated its intended operating principle. Namely, when the input flux was sufficiently high to initiate boiling, the out of plane vaporization and subsequent in-plane convection of latent heat dramatically decreased the temperature gradient across the evaporator.
- 2) Using a variety of empirical correlations relating non-visible surface temperatures to recorded IR data, the heat transfer analysis for estimating convective losses and calculating effective conductivity was greatly improved over the isothermal model presented in Chapter 7.
- 3) The veracity of this heat transfer analysis was verified experimentally by analyzing the special case of pure conduction and correctly predicting the intrinsic solid conductivity of silicon over a wide range of heat fluxes and surface temperatures.
- 4) As the input flux was increased and more vigorous vaporization occurred, several distinct performance regimes were observed. These were identifiable not only by representative conductivity values, but by their fluctuations and the phase of the ejected fluid.
- 5) In one particular performance regime, approximate periodicity was observed in the effective thermal conductivity. While not established with certainty, it is believed that either dissolution and subsequent re-solution of gas bubbles or transient phase change could cause such behavior.

- 6) After the onset of vaporization, an effective thermal conductivity of 1000-2000 W/m·K is easily maintained and very insensitive to further increases in flux. This demonstrates a very close approximation to the ideal thermal ground plane.
- 7) The highest effective thermal conductivities were always observed just prior to dryout. Single data point spikes as well as short (3-5 second) intervals of 10,000-20,000 W/m·K were regularly observed.

8.1 Solid Conduction vs. Vaporization: Concept Verification

As discussed in Section 7.2.2, before attempting any quantitative conductivity measurements, it made sense to first test the basic operating principle of the open-loop evaporator and test stand. The temperature distributions of two linear and two areal domains on the evaporator surface were analyzed during various operational modes to confirm temperature gradient and uniformity changes at the onset of vaporization (see Fig. 5.4). I first examined the devices operating at steady state conditions, without any fluid flow, as the dissipated power was incrementally increased. The four graphs in Fig. 8.1 below plot the maximum, minimum, and mean temperatures of the linear and areal domains for dry tests over a range of 1.48-3.52 W. For convenience, a condensed version of Fig. 8.14 has been included in the upper left corner of each plot, with the corresponding linear or areal domain highlighted.

Obviously, there's not much going on in these plots; one can see a nearly linear increase in temperature for all four domains, with the small domains (Line 2 and Box 2) showing less variation between maximum and minimum. Given that these domains have fewer pixels and are completely contained within the larger domains, this is guaranteed. Likewise, Box 1, which covers the entire evaporator, will always show the largest absolute variation. Lastly, it appears that the difference between maximum and minimum temperature increases slightly with power (and thus temperature). This can be seen more clearly in Fig. 8.2, which plots the average temperature gradient (in degrees K per 1000 pixels, abbreviated "kPix") along the linear domains and the standard deviation in temperature over the areal domains. With increased power, we would expect a larger temperature gradient because under a solid conduction model, a larger temperature difference is required to transfer more heat. Moreover, although the system is not *strictly* governed by solid conduction (there are convective losses as yet unaccounted for), note that doubling the power from 1.5 to 3W roughly doubles the temperature gradient of Line 1 (from 50-100 K/kPix). This is exactly what would be predicted by Fourier's Law (Eq. 8.5). We probably should not expect so close a match for the small domain, since it is only nine pixels in length and thus much more sensitive to noise. Regarding the variation in standard deviation over the areal domains, the results are also expected. An increasing in the gradient over a fixed domain will inherently increase the range, and furthermore the *absolute* increase will be larger for the larger domain.

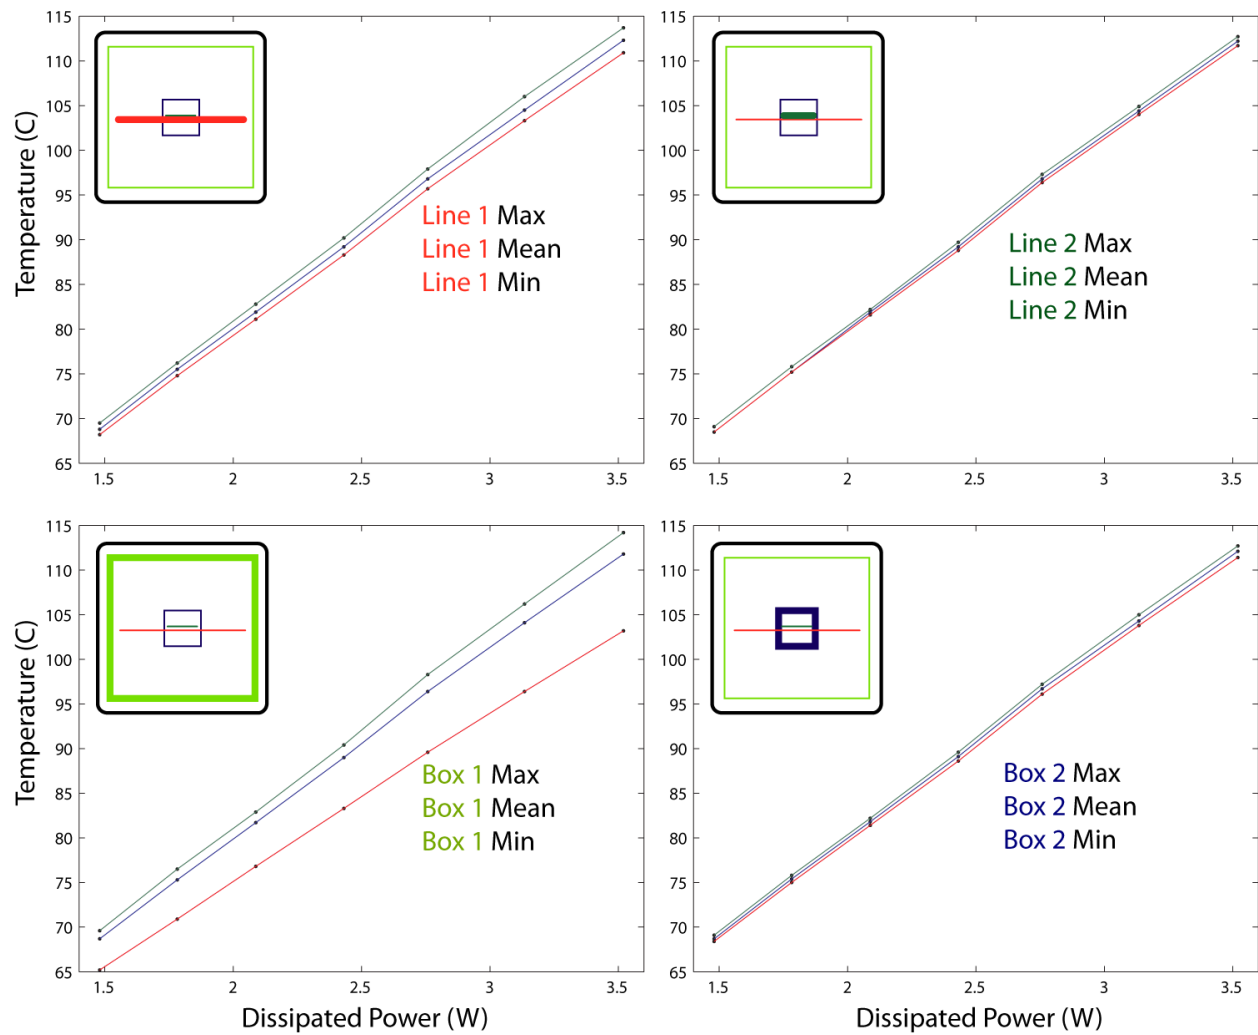


Fig. 8.1 – Variation in domain temperature statistics for increasing power in a representative dry test.

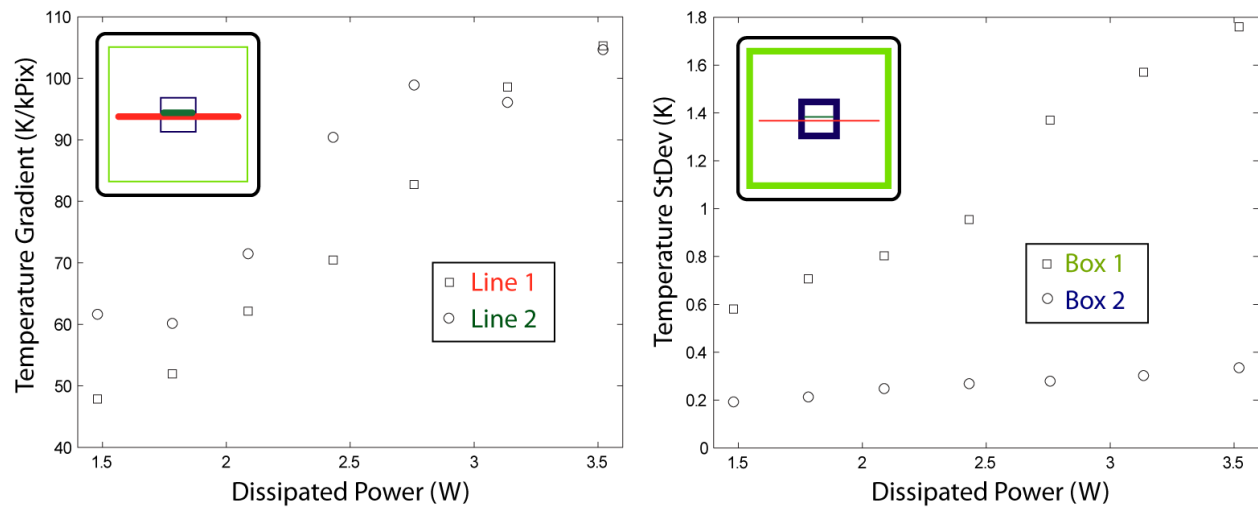


Fig. 8.2 – Temperature gradient and standard deviation for increasing power in a representative dry test.

Having qualitatively confirmed the expected heat transfer characteristics of the system under conditions largely dominated by solid conduction (and also having approached the maximum allowable dry system power dissipation that precludes melting the Delrin clamps), I next looked at the behavior of the system when fluid was pumped through the device. (Note: for all of the tests described in this section, the flow rate provided by the syringe pump was 5400 $\mu\text{L/hr.}$) For the wet tests, I continued the voltage ramp-up of the ceramic heater right where the dry tests left off (16.0V), but due to the temperature dependence of the heater resistivity, the current and thus the dissipated power at 16.0 V was slightly higher for the water-cooled system than the dry system (3.6 W vs. 3.52 W). Figures 8.3 and 8.4 depict the same data for the wet tests that Figs. 8.1 and 8.2 depicted for the dry tests (maximum, minimum, and mean temperatures of the linear and areal domains; average temperature gradient along the linear domains; and standard deviation in temperature over the areal domains) for power levels ranging from 3.6-5.34 W. Furthermore, although the absolute limits are different, the scales of the x and y axes in Fig. 8.1 and 8.3 are identical, so direct visual comparisons can be made more easily.

The first observation I wish to point out is the continued linear increase in the maximum, minimum, and mean temperatures of all the domains from 3.6 to approximately 4.5 W. Due to the initial cooling effect of the injected water *via* sensible heating, the surface temperatures at 3.6 W are significantly lower than those for the dry test at similar power ($\sim 90^\circ\text{C}$ vs. $\sim 112^\circ\text{C}$), but otherwise the behavior is quite similar. The slopes are also somewhat lower than those of the dry tests (indicating a small increase in effective thermal conductivity), but the curves' continued linearity suggest that the heat transfer is still largely dominated by solid conduction in the silicon. A second observation is that the average temperature gradient and standard deviation continue to increase nearly linearly until the power reaches 4.5-5 W. This seems to indicate that, apart from the sensible heating of the injected water, no *significant* new heat transfer phenomena are present during single phase liquid flow (see Section 8.4.2 for further discussion on the possible effects of single phase convection). Just as was the case in the dry tests, temperature gradients and areal variation increase nearly linearly as the input flux is increased, following Fourier's Law.

As the dissipated power approaches 5 W, however, we see dramatic changes in the device behavior. First of all, there is a clear decrease in the slope of the temperature curves, and in fact, they almost level off completely at the highest power levels plotted. This is, by definition, an approximation of the idealized thermal ground plane: a surface whose temperature does not change even as the heat flux is increased. Even more dramatically, both the average temperature gradient and the areal standard deviation begin to *decrease* rapidly, even as the dissipated power and surface temperature *increase*. Such behavior is in direct violation of Fourier's Law and undoubtedly indicates new transfer phenomena. That the source of this change was fluid vaporization and in-plane convection of latent heat was evidenced *visibly* by the ejection of vapor from the device outlet channels. Moreover, the decrease in temperature variation over the areal domains further reinforces the concept of a thermal ground plane. In summary, the onset of vaporization in the evaporator caused a tremendous increase in the effective thermal conductivity of the device, which resulted in a more isothermal surface. This is precisely what the device was designed to do and precisely what these tests were designed to confirm. The next three sections in this chapter will attempt to quantify this performance increase.

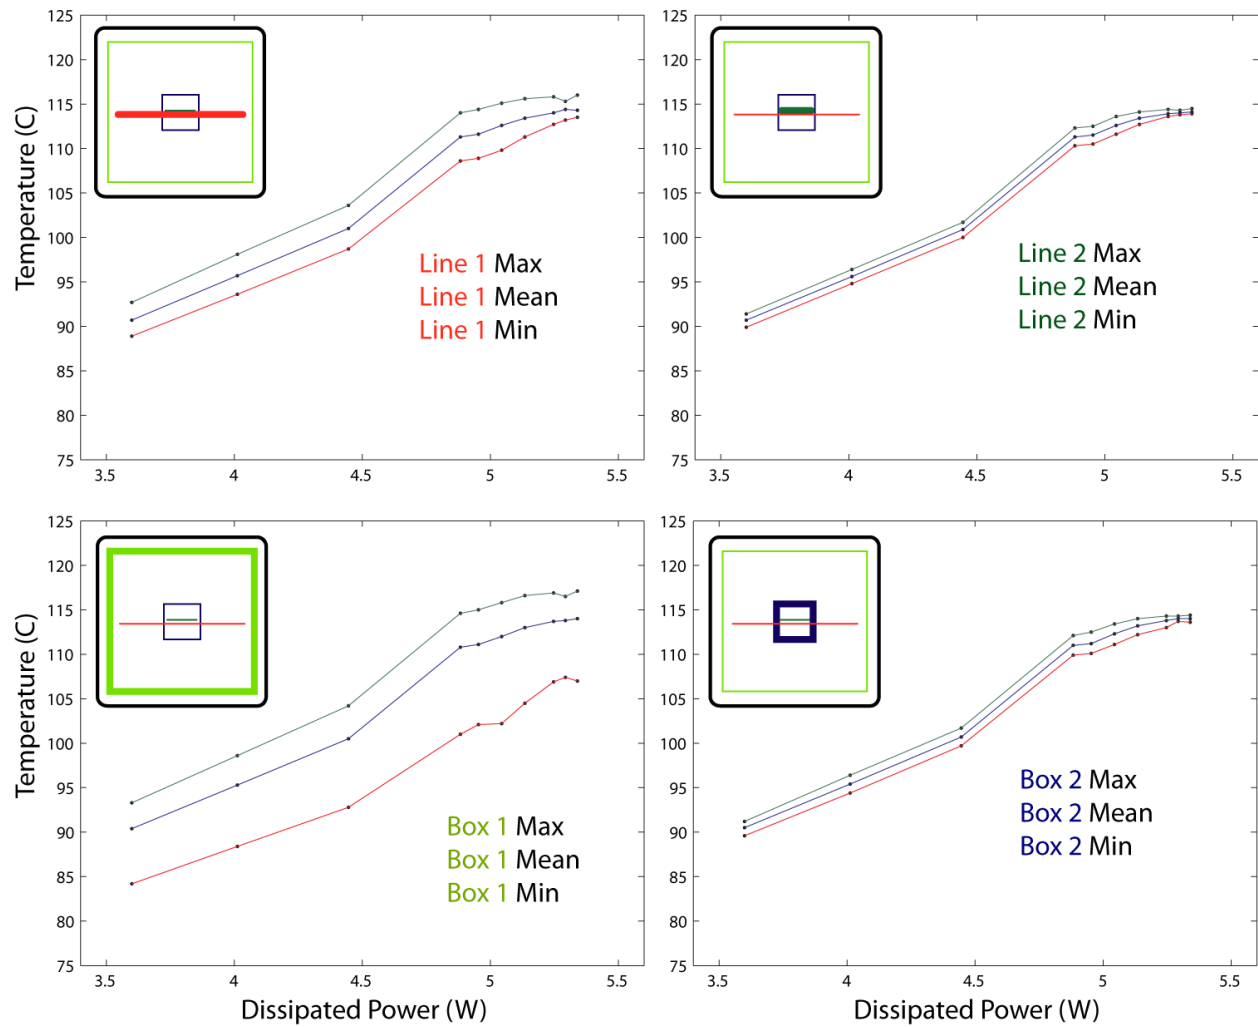


Fig. 8.3 – Variation in domain temperature statistics for increasing power in a representative wet test.




Fig. 8.4 – Temperature gradient and standard deviation for increasing power in a representative wet test.

8.2 Improving the Isothermal Clamp Model

After about a week of conductivity testing, I began to seriously reconsider my choice of the areal domain used to find the average surface temperature of the heater clamps (*i.e.* “Box 1” in all top view IR screen shots). Remember, the higher this temperature, the greater the convective losses predicted, and thus the lower the calculated effective thermal conductivity. As discussed in Section 7.5, I had initially placed the box over the footprint of the underlying ceramic heater such that, if anything, I would *overestimate* the convective losses and thus *underestimate* the effective conductivity. Though my initial motivation for conservatism remained, it became clear that the average surface temperature of the heater clamp was significantly less than the average temperature of Box 1 and that this was significantly reducing the calculated conductivity.

I have stated before that the standard deviation in temperature for the heater clamp was typically less than 5°C. In fact, the standard deviation *within* Box 1 was typically less than 2°C. However, that 5°C figure was more typical of the *top surface*—not the entire clamp. I think I initially believed that convection from the top surfaces would dominate the parasitic losses. This was based partially on my experience with free convection in everyday life (*e.g.* hot asphalt, boiling pasta, *etc.*) and partially on the fact that the top surface (directly above the ceramic heater) *is* indeed the hottest surface. Consequently, I think I focused too narrowly on obtaining a characteristic (and conservative) temperature for the top surface, rather than thoroughly considering the entire system. Furthermore, I have little doubt that this psychological bias was further exacerbated by the fact that only the top surfaces of the clamps were visible in the IR images.

The major problem with this line of reasoning, however, is that my technique for estimating convective losses assigns the Box 1 temperature to all the clamp surfaces. Now, this wouldn’t be so bad if the top surface were much larger than or even similar in size to the lateral surfaces, but for the geometry of the clamps in this study, the total lateral surface area is more than 2.5 times larger than the top surface. Since total heat transfer is a function of surface area as well as temperature and convection coefficient, the greater the difference in surface area between lateral and top surfaces, the worse the overestimation of the lateral heat loss becomes. Even if convection from the top surface *is* actually much greater than from the sides, my technique will produce the opposite effect, based purely on the difference in total area. This is exactly what happened; using Box 1 temperatures as characteristic for the entire clamp, the lateral heat loss in every test was roughly twice that of the top surface.

8.2.1 – Corrective Temperature Correlations

At this point, unfortunately, I could not go back in time and redraw Box 1 over a different area that better estimates the average temperature of the entire clamp—say, further away from the footprint of the ceramic heater and closer to the periphery of the top surface. And I certainly did not want to go back and retest all of the floor-ceiling pairs already tested. What I attempted instead was to correlate the recorded Box 1 temperature to adjusted, more characteristic temperatures for use in the convective loss calculations. As I discussed earlier, ideally I would have used multiple IR cameras to obtain side and bottom views for characterizing those surface temperatures, but that was not an option. Furthermore, it was not practical (nor prudent!) to continually detach and remount the IR camera for side and bottom views during *each* of the

conductivity tests. However, it's a relatively simple matter to incrementally ramp up the Box 1 temperature through the entire range encountered in this study while periodically recording top, bottom, and lateral IR images. Because the geometry of the clamps is fixed (and because their thermal conductivity is very low compared to the silicon), it is quite reasonable to assume that the lateral and bottom surfaces will approximately match their Box 1-correlated temperatures. In other words, if, during a *specific* steady-state conductivity test, IR images of the bottom surface indicate an average temperature of 58°C when the average temperature of Box 1 is 66°C, then it is reasonable to assume that *any time* Box 1 is 66°C, the bottom surface will be 58°C. All that is required is a function (or plot) that converts recorded Box 1 temperatures into the corresponding bottom and lateral surface temperatures.

I did not believe that this was as big of a problem with the device clamps. First of all, simply from visual inspection of the IR images, the device clamps seemed much closer to the isothermal idealization. Moreover, Box 2 (the areal domain analyzed to characterize the average device clamp temperature) was not “biased” in the same way as Box 1; only half of its footprint overlapped the hot underlying device. In addition, compared to the heater clamps, the device clamps’ surface temperatures are significantly lower, so their convective heat loss are smaller and thus less influential on the total losses and calculated conductivity. However, since I was already going to collect multi-view IR imagery and perform the corrective analysis on the heater clamps, it required minimal effort to perform the same analysis on the device clamps.

Finally, since the software allows four total domains for analysis, in addition to the original Box 1 and Box 2 data, I also collected areal statistics on regions near the periphery of the clamps in the top and bottom views. This would allow me (if desired) to break up the top and bottom surfaces and calculate weighted averages for the total heat loss from that surface. In other words, I performed the entire heat transfer analysis (*i.e.* $T_{surf} \rightarrow Ra \rightarrow Nu \rightarrow h$) for two different surface temperatures and then had the option of weighting the results by their representative areas on the surface in question (for more detail, see Section 8.2.2). Similarly, for the side view images, I was able to analyze the top and bottom clamp components separately by collecting areal data on each of the four lateral surfaces in view.

Figure 8.5 below shows a representative sample set of IR images of the top, bottom, and side of the setup at steady state, illustrating the differences between the Box 1 and Box 2 temperatures collected during the conductivity tests and the actual temperatures of the various clamp surfaces. This multi-view surface temperature data was collected over the range of Box 1 and Box 2 temperatures encountered in this study (at 1 V increments), and the results are plotted in Fig. 8.6 below. For all surfaces, the data shows a strong linear dependence over the domain of interest, and I based my correlations on a linear fit using MATLAB’s built-in polyfit function. The resulting best-fit corrective correlations for each of the clamp surfaces are listed below, with Eqs. 8.1 – 8.5 pertaining to the heater clamps (functions of the top view Box 1 temperature) and Eqs. 8.6 – 8.10 pertaining to the device clamps (functions of top view Box 2 temperature).




Fig. 8.5 – Representative set of multi-view IR images used to retroactively correlate various clamp surface temperatures to the original top view Box 1 and Box 2 areal data. (Note: the color scales of these images are not the same; visual comparison can be made only within a single image.) (Device B3B6, Heater Voltage = 16.0 V)

Correlations for heater clamp surfaces (functionally dependent on Box 1 temperature):

$$T_{HC,top,periphery} = f(T_{Box1}) = 0.7159 T_{Box1} + 8.4628 \quad (8.1)$$

$$T_{HC,bottom,center} = f(T_{Box1}) = 0.8939 T_{Box1} + 5.4790 \quad (8.2)$$

$$T_{HC,bottom,periphery} = f(T_{Box1}) = 0.5722 T_{Box1} + 16.4191 \quad (8.3)$$

$$T_{HC,lateral,top\ clamp} = f(T_{Box1}) = 0.7967 T_{Box1} + 7.9244 \quad (8.4)$$

$$T_{HC,lateral,bottom\ clamp} = f(T_{Box1}) = 0.7172 T_{Box1} + 9.5940 \quad (8.5)$$

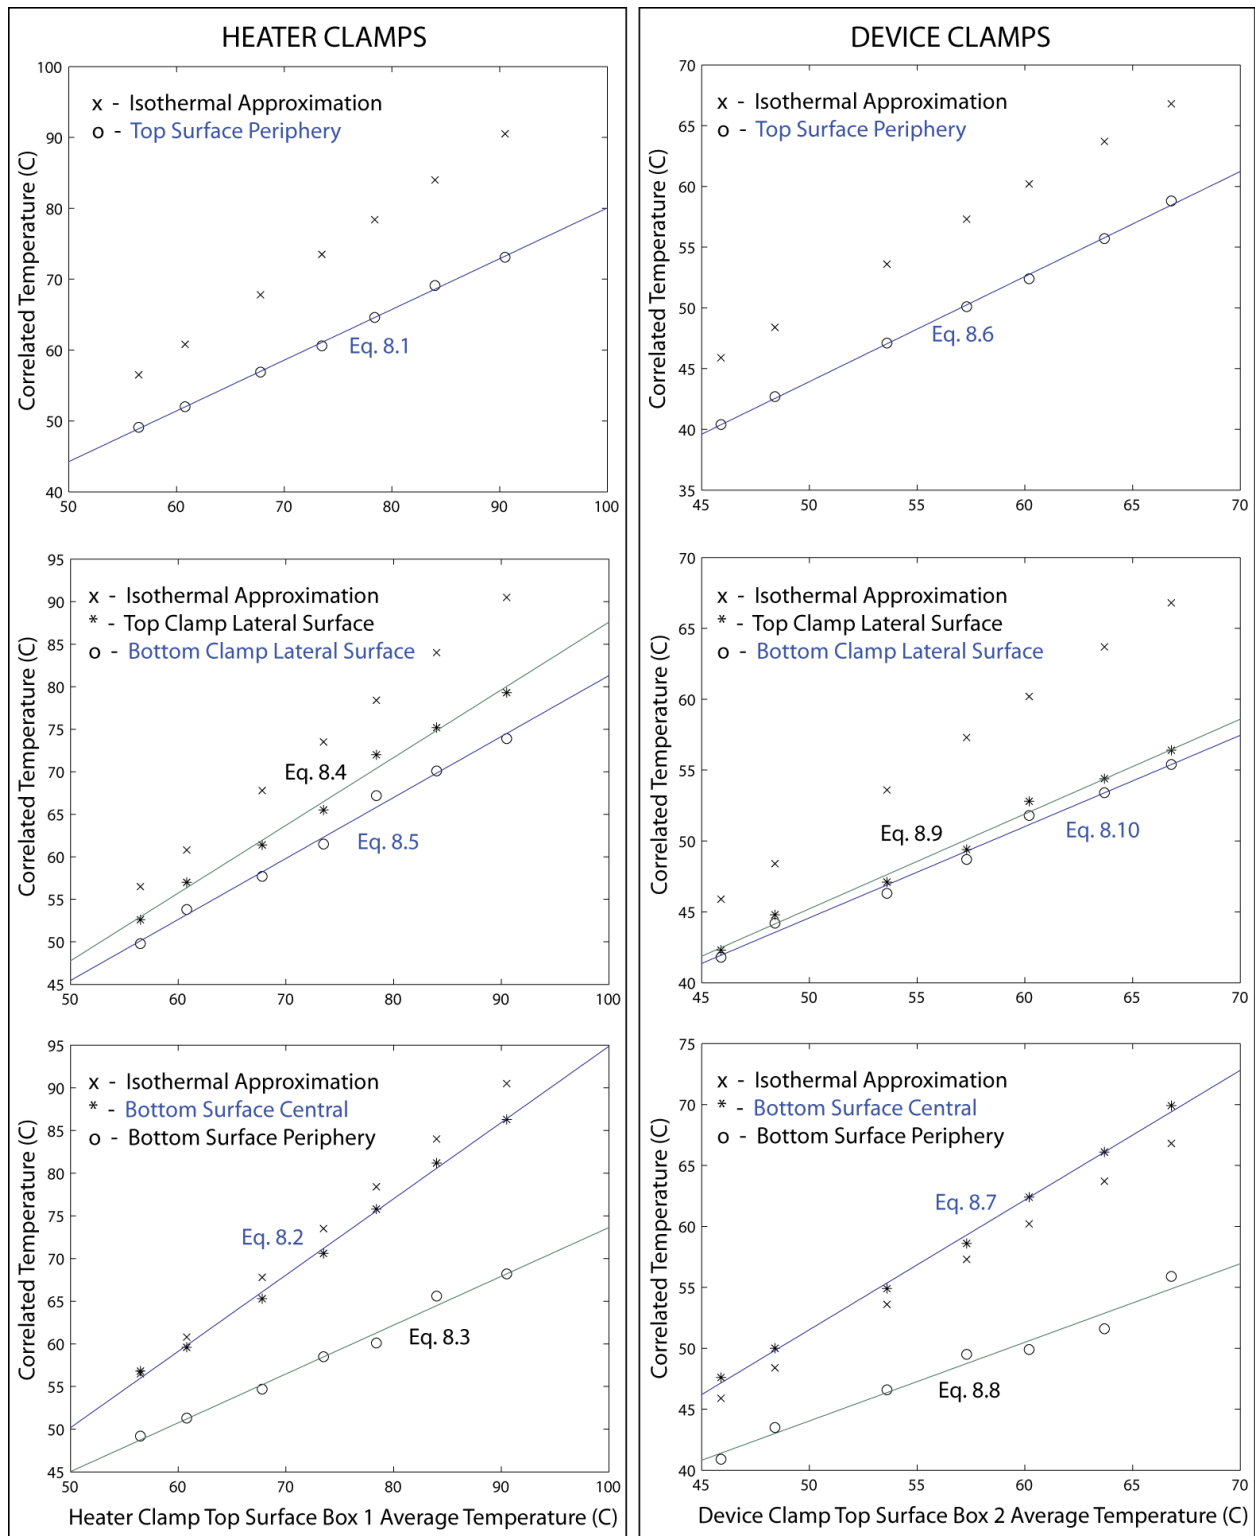


Fig. 8.6 – Multi-view temperature data plotted to correlate various clamp surface temperatures to the original top-view data of Box 1 and Box 2 (Device B3B6, Heater Voltage = 16.0 V).

Correlations for device clamp surfaces (functional dependence on Box 2 temperature):

$$T_{DC,top,periphery} = f(T_{Box2}) = 0.8647 T_{Box2} + 0.6944 \quad (8.6)$$

$$T_{DC,bottom,center} = f(T_{Box2}) = 1.0638 T_{Box2} - 1.6628 \quad (8.7)$$

$$T_{DC,bottom,periphery} = f(T_{Box2}) = 0.6446 T_{Box2} + 11.8120 \quad (8.8)$$

$$T_{DC,lateral,top\ clamp} = f(T_{Box2}) = 0.6687 T_{Box2} + 11.7830 \quad (8.8)$$

$$T_{DC,lateral,bottom\ clamp} = f(T_{Box2}) = 0.6437 T_{Box2} + 12.3960 \quad (8.10)$$

The plots in Fig. 8.6 clearly show that a Box-1-based isothermal clamp approximation will overestimate the convective losses, especially for the lateral surfaces of the clamps. Interestingly, the Box 1 and Box 2 temperatures were actually very close to the temperature of the central regions on the bottom surfaces of the clamps, suggesting that the overestimation of convective losses due to the isothermal approximation was based almost entirely on the overestimation from the lateral surfaces. In fact, the average surface temperature in the central region of the bottom surface of the device clamps was actually *higher* than the Box 2 value; thus the original isothermal approximation actually *underestimated* slightly the heat transfer coefficient (and thus convective loss) from that surface. In hindsight, this is not surprising, given that the bottom device clamp contained the (near-saturated) liquid reservoir, which was in direct thermal contact with the silicon device over the *entire area* of the evaporator. The top device clamp, on the other hand, was only in contact with a 2 mm wide border surrounding the evaporator region. Regardless, this further demonstrates just how influential the calculated lateral heat loss is with respect to the total heat loss; underestimated the heat loss from the bottom was overwhelmed by the overestimation from the sides.

The attentive reader will note that all I have done here, in essence, is to improve upon the original assumption of isothermal clamps (see Section 7.5). Rather than using just one characteristic temperature to calculate the convective losses for the entire clamp, I instead use that one recorded temperature (in this case, the average temperature in Box 1 or Box 2), along with the empirical correlations in Eqs. 8.1 – 8.10, to better estimate characteristic temperatures for *each* surface. Each of *these* temperatures is then used in the analysis detailed in Section 7.4 to find the heat lost from that particular surface. While this method is still an approximation (with its own associated “area lumping” error), there is no doubt that it represents an improvement over the isothermal clamp model.

8.2.2 – Effects of Corrective Temperature Correlations

As a concrete example, Table 8.1 below summarizes, for a representative test, the calculated heat transfer coefficients and convective losses from various surfaces of the heater clamps, allowing direct comparison of the original isothermal approximation to the correlated surface temperature model. Table 8.2 summarizes the same data for the device clamps from the same test. As discussed above, using the correlated temperatures decreases the calculated heat transfer from all of the surfaces except the bottom of the heater clamp, whose correlated temperature is slightly higher than the Box 2 average. (Note: obviously, the heat transfer from

the top surfaces is unchanged, since both models use the raw Box 1 and Box 2 temperatures for those surfaces.) However, the decreases in lateral heat transfer due to significantly lower correlated temperatures (10.5°C lower for the heater clamps; 7.9°C for the device clamps) more than make up for this increase. Thus the overall result is a decrease in the calculated convective losses. In this example, the reduction was 15% for the heater clamp, 18% for the device clamp, and 16% for the entire system. Taking into account at *all* of the wet tests performed, when using the corrective correlations, heater clamp loss decreases ranged from 14-16%, device clamp losses from 18-20%, and total system losses from 16-17%.

Table 8.1 – Heater clamp calculated values: isothermal approximation vs. correlated surface temperatures (Device B3B6, Heater Voltage = 16.0 V).

	Top Surface			Bottom Surface			Lateral Surface			Total
	T (°C)	h (W/m ² ·K)	Q_{lost} (W)	T (°C)	h (W/m ² ·K)	Q_{lost} (W)	T (°C)	h (W/m ² ·K)	Q_{lost} (W)	Q_{lost} (W)
Isothermal Approx.	79.2	13.30	0.378	79.2	6.65	0.189	79.2	11.41	0.817	1.384
Correlated Temperature	79.2	13.30	0.378	76.3	6.56	0.176	68.7	10.89	0.623	1.177

Table 8.2 – Device clamp calculated values: isothermal approximation vs. correlated surface temperatures (Device B3B6, Heater Voltage = 16.0 V).

	Top Surface			Bottom Surface			Lateral Surface			Total
	T (°C)	h (W/m ² ·K)	Q_{lost} (W)	T (°C)	h (W/m ² ·K)	Q_{lost} (W)	T (°C)	h (W/m ² ·K)	Q_{lost} (W)	Q_{lost} (W)
Isothermal Approx.	58.3	12.26	0.164	58.3	6.13	0.082	58.3	9.80	0.437	0.683
Correlated Temperature	58.3	12.26	0.164	60.4	6.22	0.089	50.4	9.22	0.306	0.559

Interestingly, these relative decreases in convective losses tend to *increase* with input heat flux. In other words, the largest difference between the isothermal calculation and the correlated temperature calculation occurs at the highest dissipated power. This makes intuitive sense. As the input heat flux increases and heat is convected more efficiently (as vapor) *through the device*, an increasingly smaller fraction of the heat must be dissipated through the clamps. Furthermore, if less heat is conducted through the clamps, then the clamp surface temperatures required to convect that same heat to the ambient will be lower. This means that at higher power, the isothermal model more grossly overestimates the surface temperatures (and thus the convective losses), while at low power (when a greater fraction of the heat must be conducted through the clamps), the temperature overestimation of the isothermal model is less significant. Continuing the example with the same device and test sequence referenced above, Fig. 8.7 below plots the loss discrepancy between the two methods over a range of dissipated powers from 3.6-5.1 W. Note: the data point analyzed in Tables 8.1 and 8.2 corresponds to the lowest power level (3.6 W) in Fig. 8.7.




Fig. 8.7 – Difference in calculated convective losses when using the isothermal approximation vs. the correlated temperature model (Device B3B6, Heater Voltage = 16.0-19.5V).

Another interesting result was that the use of the correlated temperature model did not significantly affect the *relative* heat lost by each clamp with respect to the other. In the specific example above, the heater clamp percentage increased from 67.0% to 67.8%. These results are entirely typical. In 82 wet tests that formed the bulk of the experimental data, the isothermal model predicted that the heater clamps accounted for 65-69% of the total heat lost, while the correlated temperature model predicted 65-70%. On average, the correlated temperature model increased the heater clamp percentage by 1%; in no test was the difference more than 1.2%. These numbers provide solid evidence for the fidelity of the correlations.

The primary effect of utilizing the correlated surface temperatures, of course, is a decrease in the overall calculated convective loss when compared to the isothermal approximation. This, in turn, increases the calculated thermal conductivity via Eqs. 8.6 and 8.7, restated below.

$$k = \frac{Q}{A_{cross}} \frac{\Delta x}{\Delta T} \quad (8.11)$$

$$Q = Q_{dissipated\ by\ heater} - Q_{lost\ to\ convection} \quad (8.12)$$

Before discussing the quantitative effects of the correlated temperature model upon the calculated effective thermal conductivity, I want to briefly address the use of the areal weighting factors introduced in the previous section. To quickly recap, in the development of corrective correlations for the various surfaces, additional areal temperature measurements were made at the periphery of the top and bottom surfaces (Boxes 3 and 4 in the top and bottom views, Fig. 8.5). This data was then used to develop correlations (Eqs. 8.1, 8.3, 8.6, 8.8) to estimate characteristic temperatures at the periphery of the clamp surfaces based solely on the Box 1 and Box 2 data. These temperatures were then used in the Section 7.4 analysis to calculate more accurate heat transfer coefficients for the periphery of the top and bottom surfaces.

These peripheral h values, combined with the h values based on the centrally-located Box 1 and Box 2 temperatures, provided me with upper and lower limits for each surface. But we can do even better than this. These limits can also be used in a “weighted” fashion with respect to their representative areas to estimate the total heat loss from that surface. In other words, for the case of the top surface of the heater clamp:

$$\dot{Q}_{HC,top} = h_{Box1} A_{center} (T_{Box1} - T_{ambient}) + h_{Box3} A_{periphery} (T_{Box3} - T_{ambient}) \quad (8.13)$$

If we then break up the top surface and assign a fractional percentage of the total surface area that is best characterized by each domain temperature (*i.e.* $X\%$ of the top surface is closer to the central Box 1 temperature, $Y\%$ is closer to the peripheral Box 3 temperature), we have:

$$\dot{Q}_{HC,top} = A_{top} [X h_{Box1} (T_{Box1} - T_{ambient}) + Y h_{Box3} (T_{Box3} - T_{ambient})] \quad (8.14)$$

If we set $X = 1$ and $Y = 0$, we obtain the upper limit for convective loss, while $X = 0$, $Y = 1$ corresponds to the lower limit. Analogous expressions can be written for the bottom surface of the heater clamp and for both top and bottom surfaces of the device clamp. The issue then becomes how to choose appropriate weighting factors X and Y .

Obviously, in an actual operating device, there is not a sudden discontinuity in temperature separating the central and peripheral regions, but instead a continuous temperature distribution from center to edge. This is evident in the top two IR images in Fig. 8.8, which depict devices operating near the two power extremes explored in this study. The absolute temperatures are, of course, quite different (note the different color scales), but the smooth temperature gradient across the clamps is similar. How should such surfaces be divided? Given the largely linear dependence of convection upon surface temperature (at least over small temperature ranges), a natural approach would be to plot an isotherm corresponding to the *average* of the data-based central temperature (Boxes 1 and 2 for the heater clamp and device clamp, respectively) and the correlated peripheral temperature (Boxes 3 and 4). In other words:

$$T_{HC,mean} = \frac{1}{2} (T_{Box1} + T_{Box3}) \quad (8.15)$$

$$T_{DC,mean} = \frac{1}{2} (T_{Box2} + T_{Box4}) \quad (8.16)$$

Such isotherms divide each of the clamp surfaces into two distinct regions:

- 1) Surface temperatures above $T_{mean} \rightarrow$ designated as “central”
- 2) Surface temperatures below $T_{mean} \rightarrow$ designated as “peripheral”

Two IR images are shown below in Fig. 8.8. One depicts the system dissipating 2.09 W (labeled “low power”), while the other depicts the system dissipating 5.31 W (labeled “high power”). Beneath each IR image, the two regions (central and peripheral) have been color-mapped, with a unique 2-color map corresponding to the specific mean isotherms of the heater clamp and device clamp at both low and high power (four maps total). In all four maps, central regions are shaded yellow, while peripheral regions are shaded red. Uncolored regions have temperatures below the minimum temperature of the particular clamp being analyzed and have been omitted for clarity.

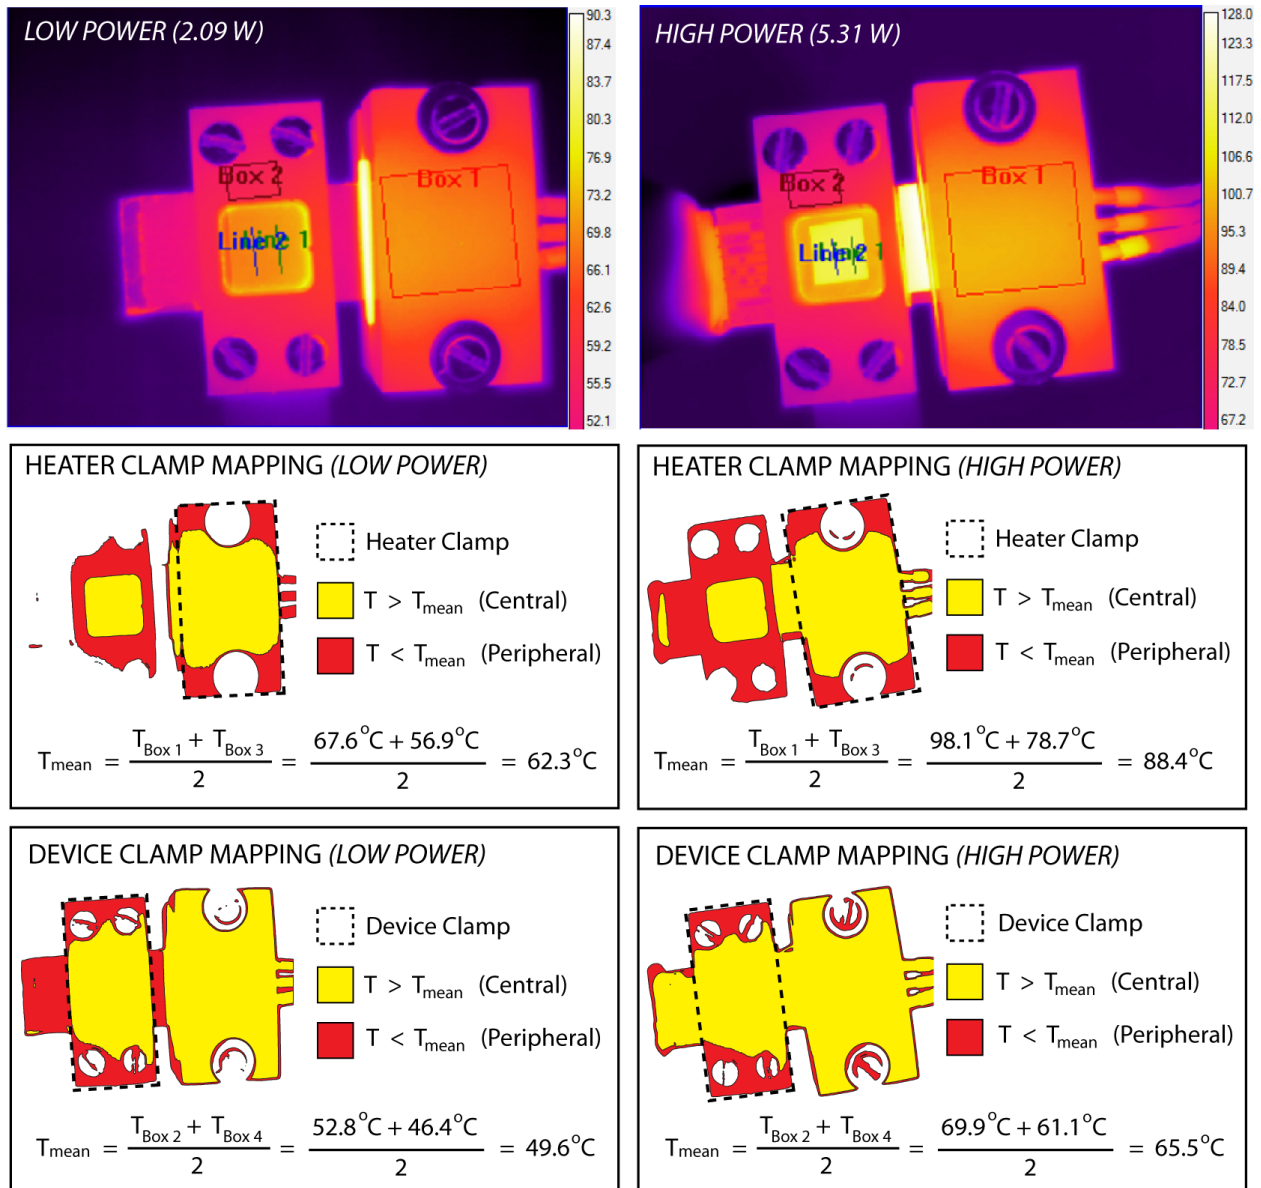


Fig. 8.8 – IR images of a device operating at low and high power, along with two-tone color maps divided along the isotherm corresponding to the mean temperature between the data-based central region and the correlation-based peripheral region of the particular analyzed clamp.

First, note that the bolts and washers are largely uncolored. This does *not* mean that their temperature is below the lowest clamp temperature. As mentioned in Section 7.5, the IR camera was calibrated more closely to the emissivity of the clamps and Kapton tape, and thus metal surfaces (with much lower emissivity) appear artificially “cold”. However, given the high thermal conductivity of steel compared to the clamps, it is reasonable to assume that the omitted regions would more or less follow the background isotherms, although such an assumption is not essential to the following analysis. In any case, it is extremely unlikely that the metallic surfaces are significantly hotter or colder than the underlying clamps.

This brings us back to the question of determining the weighting factors X and Y . The key areas to focus upon in Fig. 8.8 are the rectangular regions enclosed by the dotted lines; these indicate the outline of the particular clamp being analyzed. Once again, yellow areas are closer in temperature to the original Box 1 or Box 2 data and are thus included in the central region; red areas are closer to the correlated temperatures of Box 3 or Box 4 and are thus included in the peripheral region. How much of the total surface area does each of these irregularly-shaped regions represent? Figure 8.9 below shows a range of *regularly*-shaped areal partitions for a properly-proportioned heater clamp. Compare the white/gray divisions of the rectangles to the yellow/red divisions enclosed within the dotted rectangles. Without necessitating precise areal calculations, it is clear that in all four images, the peripheral region is certainly larger than 10% or even 20% of the total area. On the other hand, assigning 50% to the periphery is probably too generous. Some of the images are closer to 30%, others to 40%—perhaps 35% is best—an exact answer is not required. After all, this analysis was performed merely to *improve* upon our estimate with respect to the upper and lower limits established earlier.

Keep in mind that the two images analyzed in Fig. 8.8 represent the upper and lower limits of dissipated power. Furthermore, the low-power data comes from a “dry” baseline test, while the high-power data comes from a fully vaporizing “wet” test. *This is significant.* If the central/peripheral divisions present under widely diverse tests conditions are both approximately 65/35 (*i.e.* $X = 0.65$, $Y = 0.35$), it is probably reasonable to assume that such divisions are characteristic of *all* of the tests performed at intermediate power levels. Lastly, it should be noted that identical analysis was performed on the bottom surfaces at both ends of the power spectrum, and similar central-peripheral divisions were observed. Thus it appears that Eq. 8.14 (even with constant weighting factors!) should provide a legitimate improvement over a correlated temperature model that assumes a constant (but unique) surface temperature for each clamp surface.

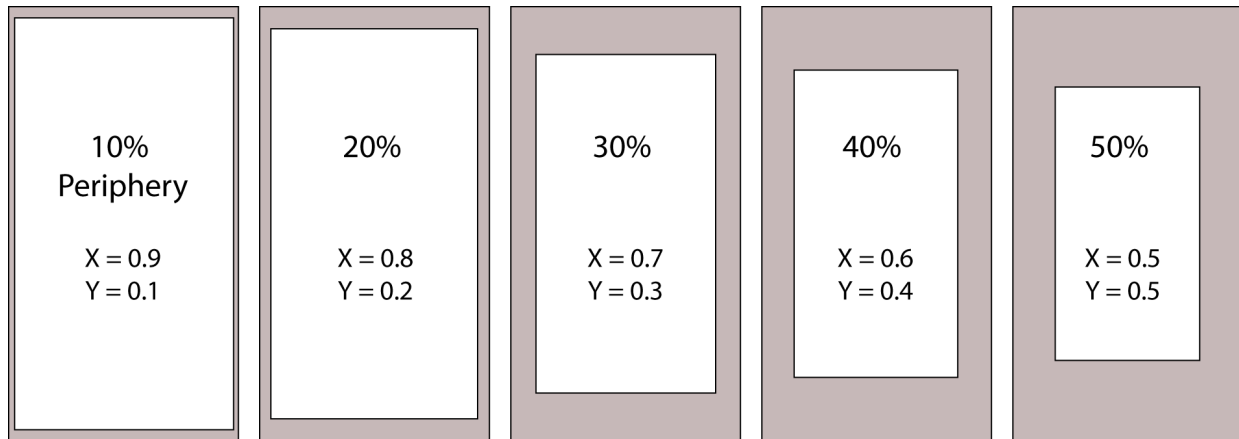


Fig. 8.9 – Graphical view illustrating a range of central/peripheral partitions for a properly-proportioned heater clamp top surface. White is central; gray is peripheral.

Setting aside for the moment our best-guess X and Y values, let us examine the combined effects of the correlated temperature model and weighting factors on the overall convective heat loss and, more importantly, the effective thermal conductivity. As stated near the beginning of this section, correcting for the overestimation in surface temperature will no doubt cause the

calculated convective losses to decrease and the effective thermal conductivity to increase, *but by how much?* Tables 8.3 and 8.4 summarize the calculated heat losses and effective conductivities for the two examples considered above (*i.e.* the devices in Fig. 8.8). Central weighting factors of 0.5 to 1.0 are considered, with our best-guess of 65/35 indicated in bold. In addition, the results of the original isothermal model are given at the bottom of each table for comparison. I have also plotted a continuous distribution of the effective thermal conductivity as a function of central weighting factor X for both examples (see Figs. 8.10 and 8.11).

Table 8.3 – Effects of central/peripheral weighting factors on convective losses and effective thermal conductivity (Device B2C4, Heater Voltage = 12.0V, Power = 2.09 W, Dry).

Central Weight X	Heater Clamps		Device Clamps		System	k_{EFF} (W/m·K)
	$Q_{lost,top}$ (W)	$Q_{lost,bottom}$ (W)	$Q_{lost,top}$ (W)	$Q_{lost,bottom}$ (W)	$Q_{lost,total}$ (W)	
1.0	0.277	0.132	0.129	0.070	1.326	117.9
0.9	0.269	0.127	0.125	0.067	1.307	120.9
0.8	0.260	0.123	0.121	0.065	1.287	123.9
0.7	0.251	0.119	0.117	0.062	1.268	127.0
0.65	0.247	0.116	0.116	0.061	1.258	128.5
0.6	0.242	0.114	0.114	0.059	1.248	130.0
0.5	0.234	0.110	0.110	0.057	1.228	133.0
Isothermal	0.277	0.139	0.129	0.065	1.559	94.5

Table 8.4 – Effects of central/peripheral weighting factors on convective losses and effective thermal conductivity (Device B8C3, Heater Voltage = 19.9V, Power = 5.31 W).

Central Weight X	Heater Clamps		Device Clamps		System	k_{EFF} (W/m·K)
	$Q_{lost,top}$ (W)	$Q_{lost,bottom}$ (W)	$Q_{lost,top}$ (W)	$Q_{lost,bottom}$ (W)	$Q_{lost,total}$ (W)	
1.0	0.551	0.252	0.242	0.131	2.497	1676
0.9	0.533	0.243	0.236	0.125	2.459	1699
0.8	0.516	0.234	0.230	0.120	2.420	1722
0.7	0.498	0.225	0.224	0.115	2.382	1745
0.65	0.489	0.220	0.221	0.112	2.363	1756
0.6	0.480	0.215	0.218	0.109	2.344	1767
0.5	0.462	0.206	0.212	0.104	2.305	1790
Isothermal	0.551	0.276	0.242	0.121	3.015	1368

As expected, in addition to being linear, the correlated temperature model produces lower system heat losses and higher conductivities. The one exception is the bottom surface of the device clamps (at 0% peripheral weighting, $X = 1$), which, as discussed earlier, actually has a higher correlated temperature than the Box 2 approximation. Also, the top surfaces at 0% weighting are, of course, identical to the isothermal model, since 0% weighting *implies* a uniform Box 1 or Box 2 temperature on the top surface. I have not included the lateral surfaces of the clamps because they are unaffected by weighting, which was done only on the top and bottom surfaces. The lateral surfaces of the top and bottom of each clamp was assigned its own unique domain and temperature (See Fig. 8.5). However, one can easily calculate the lateral heat loss in each case simply by subtracting the top and bottom losses from the total system loss.




Fig. 8.10 – Effective thermal conductivity vs. central weighting factor for a device at low power (Table 8.3 example).




Fig. 8.11 – Effective thermal conductivity vs. central weighting factor for a device at high power (Table 8.4 example).

Speaking more quantitatively, the correlated (but unweighted) model predicts a total system heat loss that is 14.9% less than the isothermal model for the low-power device and 17.2% less for the high power device. This translates into conductivities that are 24.8% and 22.5% higher than the isothermal model. If we instead use our best-guess 65/35 weighting, the effective conductivities increase by 9% and 4.8%, respectively, over the unweighted values (or 36% and 28.3% over the isothermal values). While certainly significant, these increase do not suggest some entirely unexplained phenomena, as might be required if the conductivity, say, doubled or tripled. Moreover, it is certainly noteworthy (and extremely encouraging!) that the 65/35 weighting produces an effective thermal conductivity of 128.5 W/m-K for the dry baseline test—tantalizingly close to the intrinsic conductivity of silicon (130 W/m-K). In fact, all of the weighting values considered here produce dry conductivities within 10% of 130 W/m-K, and the more reasonable weighting range of 30-40% periphery produce conductivities within 3%. Whether this is simply a fluke or instead provides solid validation of the methodology and analysis presented will be the subject of the next section.

One additional tool that this analysis provides is a *single* simple correlation to estimate convective losses from each clamp in future tests based solely on the Box 1 and Box 2 temperature data (rather than requiring the full analysis from Section 7.4). To obtain this correlation, I plotted the convective losses of each clamp (using the 65/35 weighting factor) for 103 previously performed tests, as a function of their recorded Box 1 and Box 2 temperatures. This plot is shown in Fig. 8.12. Although the dependence is overwhelmingly linear ($R^2 > 0.999$), an essentially exact fit ($R^2 = 1$) can be found for the relevant temperature domains with a second order polynomial using MATLAB's `polyfit` function (Eqs. 8.17 and 8.18). This greatly simplifies real-time conductivity measurements when implementing user-defined functions in the ExaminIR software.

$$Q_{lost,HC} = f(T_{Box1}) = 4.743 \times 10^{-5} T_{Box1}^2 + 0.0173 T_{Box1} - 0.5520 \quad (8.17)$$

$$Q_{lost,DC} = f(T_{Box2}) = 6.022 \times 10^{-5} T_{Box2}^2 + 0.0123 T_{Box2} - 0.3933 \quad (8.18)$$




Fig. 8.12 – Simple polynomial correlations to estimate convective loss based solely on recorded temperatures from 103 previously performed tests.

8.3 Experimental Verification of Heat Transfer Analysis

Although the results of the previous section suggest a very good model for estimating convective losses, that model must be tested. As discussed in the previous chapter, “dry” baseline tests were performed primarily to provide some basis of comparison for subsequent “wet” vaporization tests, based on the understanding that each floor-ceiling device pair may see a different inherent “front-end” thermal resistance based on varying clamp pressure, contact resistance, *etc.* for that set of experiments. However, these tests also provide valuable experimental verification concerning the analytical techniques used to estimate the parasitic convective losses in the system. After all, if the data conversion techniques detailed in Section 7.4 and improved upon in Section 8.2 are largely erroneous, the effective thermal conductivity for a dry test could be dramatically different from the intrinsic solid conductivity of single crystal silicon ($\sim 130 \text{ W/m}\cdot\text{K}$). If, on the other hand, these methods *consistently* indicate dry device conductivities near $130 \text{ W/m}\cdot\text{K}$ (regardless of the exact setup or input heat flux), we can be ever more confident that the heat transfer analysis is sound. Perhaps the easiest way to verify the method is to simply observe the change (or lack thereof) in the calculated thermal conductivity of a dry device as the input heat flux is varied. After all, it is always possible that at *one specific* power level, a completely erroneous method will (completely by chance) calculate a thermal conductivity near $130 \text{ W/m}\cdot\text{K}$. However, if that conductivity stays relatively constant over a large range of input fluxes and surface temperatures, we gain confidence in the methods used.

Figures 8.13 and 8.14 below show the results of just such a test. Here, a single device pair was tested with no fluid flow over a range of dissipated powers from approximately 1.5-3.5 W. Though the various surface temperatures vary significantly as the input heat flux is increased, the calculated effective thermal conductivity stays relatively constant and, indeed, very close to the intrinsic value of $130 \text{ W/m}\cdot\text{K}$. In any case, it certainly does not appear to

increase or decrease in any functional way, which would be a sure sign that the convective loss analysis was faulty. For this particular test involving seven power levels, the average k_{EFF} was 132.4 W/m·K with a standard deviation of 8.88 W/m·K. Although this extended, variable-flux baseline test was not performed for all the device pairs, several pairs that were tested yielded similar results, verifying the methodology of the convective loss calculations.




Fig. 8.13 – Variation in clamp and device surface temperature during dry testing to confirm convective loss calculation methods (Device B2C4, Heater Voltage = 10.0-16.0V,




Fig. 8.14 – Corresponding variation in calculated effective thermal conductivity during dry testing to confirm convective loss calculation methods (Device B2C4, Heater Voltage = 10.0-16.0V, Dry).

8.4 General Performance Regimes

One of the first observations made when carrying out the various tests described in Sections 8.3.2 and 8.3.3 was the existence of distinct performance “regimes”, based almost

entirely on the input heat flux (and differences in mass flow, where applicable). This was, of course, entire expected, given that the very concept of the open-loop test platform was to observe changes in the temperature gradient across the evaporator as the dominant steady-state heat transfer shifts from in-plane solid conduction to out-of-plane phase-change convection (see Fig. 5.4). Moreover, as explained in Section 2.4, it is well understood that such transitions to and from boiling cause phenomenological changes in the heat transfer modes that can dramatically affect overall energy transfer.

8.4.1 – Regime I: Dry Baseline

The first distinct performance regime (which I will sometimes refer to as “dry baseline”) is associated with pure solid conduction, that is, when the syringe pump is disconnected and there is no fluid flow into the evaporator. Figure 8.15 below is a representative temporal plot of effective thermal conductivity for a dry baseline test at steady state. In this case, the data spans almost 14 minutes, and the integrated, time-average value of k_{EFF} is 129.8 W/m·K, with a standard deviation of 2.1 W/m·K. As would be expected, in addition to closely matching the intrinsic conductivity of silicon (130 W/m·K), the conductivity in dry tests is somewhat steadier than tests involving single-phase fluid flow, and *much* steadier than tests involving phase change (see Figs. 8.16-8.18 for comparison). However, the *calculated* conductivity is still very sensitive to IR temperature measurement, so some level of noise is expected, even with pure solid conduction. It should also be mentioned that even during a dry test, there is still fluid inside the device—namely air. This air will, of course, contribute to the total conductivity of the device. However, with the understanding that: a) any bulk motion of the internal air is due to natural convection (not forced airflow) and b) the thermal conductivity of quiescent air is almost four orders of magnitude less than silicon ($k \sim 30 \times 10^{-3}$ W/m·K), it is very reasonable to neglect this contribution entirely, and I have done so throughout this study.




Fig. 8.15 – Regime I: Representative temporal plot of effective thermal conductivity for a dry baseline test (Device B8C5, Heater Voltage = 14.0V, Dry).

8.4.2 – Regime II: Steady Single-Phase

The next distinctive performance regime (which I will sometimes refer to as “steady single-phase”) occurs once the syringe pump is connected and liquid is pumped through the system, but where the power level remains low enough to completely preclude phase change. Liquid outflow at the vapor channel outlets is smooth and steady. This performance change is believed to be caused by three effects. First, some of the input heat is now absorbed by the incoming liquid as it is heated to saturation (either in the reservoir or in the device itself); this has the effect of increasing the perceived device conductivity. For example, for a flow rate of 4000 $\mu\text{L/hr}$ ($\sim 10^{-9} \text{ m}^3/\text{s}$), the total power required to heat the water can be estimated as follows:

$$\dot{Q} = \dot{m} c \Delta T \Rightarrow \frac{\dot{Q}}{\Delta T} = \dot{m} c = (10^{-9} \text{ m}^3/\text{s})(1000 \text{ kg/m}^3)(4186 \text{ J/kg} \cdot \text{K}) = 0.0042 \text{ W/K}$$

For the *maximum* sensible heating required in this experiment (from 27°C to 100°C), this amounts to approximately 0.3 W.

The second effect on performance is a result of the forced convection of the flowing liquid. However, this effect is almost negligible for two reasons. First, the liquid velocities in the device are extremely low. Again, given a flow rate of 4000 $\mu\text{L/hr}$ and evaporator/channel cross-sections on the order of 10^{-6} m^2 , the characteristic liquid velocities are less than 1 mm/s. This translates to exceedingly low Reynolds numbers and thus weak convection. Second, the liquid reservoir is in direct thermal contact with the bottom surface of the silicon device, so at steady state, the inlet condition of the water is very close to saturation already. This decreases the ΔT in Eq. 2.3 to almost zero, thus minimizing convective transfer regardless of the flow velocity.

This brings me to the third effect. Even if the flow velocity was exactly zero (and forced convection was reduced to conduction into a quasi-static fluid), we would still expect a slight increase in performance simply by replacing the air in the device with water. Why? As discussed above in the section on the dry baseline regime, the conductivity of the fluid itself contributes to the overall device conductivity. Though still obviously dominated by the solid conduction of the silicon, the thermal conductivity of saturated water is approximately twenty times higher than that of air, so one would expect a *nominal* increase in the effective thermal conductivity of the entire system due simply to the introduction of working fluid.

Figure 8.16 below is a representative temporal plot of effective thermal conductivity for a steady single-phase test. In this plot, the data spans more than eight minutes, and the integrated, time-average value of k_{EFF} is 158.4 W/m·K, with a standard deviation of 7.1 W/m·K. This slight increase in effective thermal conductivity fits well with the expected results, based on the arguments above. Furthermore, as mentioned earlier, this regime is less steady than typical dry tests, but much steadier than tests involving phase change.




Fig. 8.16 – Regime II: Representative temporal plot of effective thermal conductivity for a wet test with no vaporization (Device B8C4, Heater Voltage = 19.8V).

I'd like to make one final comment regarding the sensible heat absorption of the liquid. As calculated above, the maximum sensible heat absorption is approximately 0.3 W. Coincidentally, numbers of this magnitude correspond very closely with the convective losses from individual surfaces of the clamps (see Appendix D, Step 8). Since I had come to understand very well what effect convective losses had had on the calculated thermal conductivity, I was led to wonder what *quantitative* effect the absorption of 0.3 W would have. In effect, the absorption of 0.3 W into the water is much like an extra 0.3 W being conducted through the system *without additional heating of the silicon or clamps* (i.e. having no effect on the temperature gradient or convective losses). Out of curiosity, I added 0.3 W to $Q_{conducted}$ in Step 10 of Appendix D. Upon doing so, the calculated effective conductivity increased by 45 W/m·K from 390 to 435 W/m·K. Comparing the time-average k_{EFF} values from the representative dry baseline and steady single-phase tests presented in this section, we find an increase from 129.8 to 158.4 W/m·K, or roughly 30 W/m·K. Admittedly, this is in no way a rigorous demonstration of the *mechanism* of the performance increase; after all, these two values come from entirely different tests using different floor-ceiling device pairs. However, given that most of the tests demonstrated a 20-60 W/m·K performance increase with the introduction of (non-vaporizing) working fluid, it does suggest that such an increase is compatible with the sensible heating of the inlet fluid.

8.4.3 – Regime III: Transition

The next distinct performance regime (which I will refer to as “transition”) occurs when the evaporator superheat is sufficient to substantially alter the fluid flow, which, in turn, causes a significant increase in the effective conductivity. Unlike Regime II, the liquid ejected at the

vapor channel outlets was unsteady and instead tended to “pulse” at a somewhat regular frequency (see Section 8.5 for further discussion). Two phenomena are proposed for this unsteady flow and the associated performance increase. First, it is possible that no phase change is occurring, but that dissolved gases (*i.e.* air) are simply coming out of solution (since solubility decreases with temperature). These air bubbles then grow quickly as the enclosed gas is heated in the evaporator, forcing liquid through the system at increased velocity. This rapid increase in the liquid velocity means that forced convection may no longer be negligible (as it was in Regime II) and that significant heat may be convected laterally as liquid is pushed through the device by expanding air bubbles. Furthermore, as the bubbles are carried away from the heater and cool, the gas bubbles are reabsorbed by the liquid, thus explaining the observation of pulsed fluid ejection.

A second possibility is that some phase change *is* occurring in the evaporator, but the vapor quickly recondenses in the vapor channels before reaching the vapor outlets. This transitory vaporization would also explain the pulsed flow and ejection of only liquid (no vapor). Figure 8.17 below shows a representative temporal plot of the effective thermal conductivity for this transition regime. The data spans approximately 100 seconds, and the integrated, time-average value of k_{EFF} is 410.1 W/m·K, with a standard deviation of 47.2 W/m·K. The much larger σ associated with Regime III can be explained by the unsteady nature of the flow, whether that unsteadiness is due to transitory vaporization or simply gas dissolution and re-solution. I believe it likely that the peaks in conductivity are associated with bubble growth events (either gas dissolution or vaporization), while the troughs are associated with bubble collapse (either gas re-solution or vapor condensation). I should also point out that in many tests, Regime III was very difficult to maintain for more than 10 or 15 seconds, and the device would instead quickly shift back to Regime II or progress to Regime IV (see next subsection). While at first glance, this instability may appear to be a cause for concern, it is important to remember the intended application and, more importantly, what actually represents a danger to the system being cooled. A shift back to Regime II (*i.e.* a decrease in k_{EFF}) means that vaporization ceases and thus accompanies a *decrease* in the input flux. While this may give rise to an *instantaneous* increase in surface temperature, this would also represent a power level at the extreme *lower* end of the intended operating limits and would thus be extremely unlikely to cause damage. On the other hand, a sudden progression to Regime IV (*i.e.* an increase in k_{EFF}) will instead give rise to a *decrease* in surface temperature and thus pose no danger.




Fig. 8.17 – Regime III: Representative temporal plot of effective thermal conductivity for a wet test with unsteady flow conditions (Device B2C3, Heater Voltage = 19.7V).

8.4.4 – Regime IV: Full Vaporization

The final distinct performance regime (which I will refer to as “full vaporization”) is characterized by a sudden, dramatic increase in the effective thermal conductivity and is easily recognized by the periodic ejection of vapor from the vapor channel outlets. This is the only regime where vapor is convected all the way to the condenser and thus demonstrates the operation of the device as intended. Liquid water is also ejected, so the operation is certainly not ideal, but it represents a clear phenomenological and performance departure from Regime III. Figure 8.18 below shows a representative temporal plot of the effective thermal conductivity for the full vaporization regime (after $t \approx 60$ s). The Regime IV data spans approximately 3.3 minutes, and the integrated, time-average value of k_{EFF} is 1454.1 W/m·K, with a standard deviation of 183.5 W/m·K. At no time does the effective conductivity drop below 1000 W/m·K.




Fig. 8.18 – Regime IV: Representative temporal plot of effective thermal conductivity for wet test with full vaporization, after $t \approx 60$ s (Device B3B6, Heater Voltage = 19.0V).

Figure 8.18 also illustrates a rapid, essentially monotonic progression through Regimes II and III. This does not always occur. Figure 8.19 shows the temporal plot of a device that was rapidly heated and allowed to progress through all three wet regimes. In this case, Regime III

lasted more than 50 seconds, and the increase in performance was not monotonic. In fact, at $t \approx 368$ seconds the device (performance-wise) basically reverts back to Regime II for a few seconds before reestablishing the transition regime. This figure also illustrates the rapid drop in conductivity associated with the transition to dryout after the critical heat flux has been exceeded.




Fig. 8.19 – Temporal plot of effective thermal conductivity for a device heated rapidly through all three wet performance regimes, as well as the transition to dryout (Device B8C3, Peak Heater Voltage = 20.0V).

8.5 Periodicity

One interesting and unexpected observation was the appearance of an approximate periodicity in some of the conductivity tests. Although I did not perform any kind of formal Fourier analysis, many of the temporal plots appeared to display a characteristic frequency in the rise and fall of the effective thermal conductivity, at least to the naked eye. This effect was most noticeable in tests within Regime III, that is, during the transition from single-phase liquid flow to full vaporization. For example, in Fig. 8.20, there is a sharp drop in the conductivity approximately every 4 seconds. Two additional examples in Fig. 8.21 display characteristic periods of 11 and 12 seconds; all three of these examples were in Regime III tests.

The cause of this periodicity is not known for certain, but it seems reasonable to associate the conductivity fluctuations with the pulsing nature of the liquid flow in Regime III. Though I did not formally record any data on the exact frequency of the ejected liquid, the typical period was definitely on the order of several seconds. Whether this pulsing is due to dissolution and re-

solution of gas bubbles or transitory phase change is not known (again, see Section 8.4.3 for further discussion). One thing that *is* known, however, is that these fluctuations are *not* due to the finite steps of the servo-controlled syringe pump. Not only are such pulses smoothed out by viscous effects in the microfluidic interconnect, but at the flow rate used in each of the examples shown, the stepper frequency was greater than 1 Hz. Also, it is clear that the periodicity is not strict. At times, the plots exhibit no apparent frequency at all, and sometimes several different frequencies may be observed within the same test.




Fig. 8.20 – Representative temporal plot of effective thermal conductivity illustrating approximate periodicity (Device B2C3, Heater Voltage = 19.7V).




Fig. 8.21 – Additional test examples illustrating approximate periodicity (**Left:** Device B8C4, Heater Voltage = 20.5V; **Right:** Device B3C6, Heater Voltage = 19.0V).

8.6 Maximum Performance and Design Optimization

The most significant results pertaining to maximum device performance are summarized below. These were based on tests using a very limited number of floor and ceiling combinations, but a larger range of mass flow rates and input fluxes.

- 1) After the onset of vaporization, an effective thermal conductivity of 1000-2000 W/m·K is very easily maintained and very insensitive to further increases in flux. This represents a stable boiling regime well short of the critical heat flux and likely indicates that only a fraction of the working fluid is being vaporized.
- 2) This stable boiling regime resulted in evaporator surface temperatures of approximately 113-116°C. Doubling the flux (from approximately 5 to 10 W) produced very little variation in this evaporator surface temperature. Moreover, this particular surface temperature is also consistent with the typical superheat seen experimentally for the steady vaporization of water (see Section 2.3).
- 3) After the device entered Regime III, the surface temperature of the device clamps stayed almost completely constant with increased flux, even as the total dissipated power was *doubled* (from approximately 5 W to 10 W). Conversely, the heater clamp surfaces continued to increase approximately linearly with power.
- 4) This “stable” regime was also generally insensitive to the particular device layout tested. More specifically, the variation of the conductivity *within* any particular test, effectively the “noise” produced by an inherently unsteady vaporization process, was larger than any substantive performance variation *between* devices. Let me reemphasize that I only tested a very limited number of floor-ceiling combinations, but these did cover most of the range of CAP and CFR numbers.
- 5) The highest effective thermal conductivities were always observed just prior to dryout, which was characterized by an extremely sharp drop in conductivity and a rapid rise in surface temperature. Single data point spikes as well as short (3-5 second) intervals of 10,000-20,000 W/m·K were regularly observed. Evaporator surface temperatures were approximately 117-118°C and extremely uniform (indicated by very low areal standard deviations, often less than 0.5K over the entire evaporator).

Points 2 and 3 are especially significant. These observations not only confirm the operational principle of the device, but also imply that nearly all of the excess heat (beyond that required for vaporization) is convected in-plane with little to no increase in the surface temperature. If this were not the case, *some* heat would be conducted through the device clamps and raise their surface temperature. Instead, this heat must be going purely into vaporization of the working fluid, reconfirming the supposition that throughout this regime, a fraction of the mass flow remains liquid. Most importantly, this demonstrates a very close approximation to the ideal thermal ground plane.

As stated earlier, fluctuations in effective thermal conductivity continually increase with increased conductivity, making unequivocal comparisons at peak performance difficult.

Nevertheless, it still appeared that the maximum observed conductivities were proportionally larger at higher heat fluxes and mass flow rates. In other words, although the near-dryout values were of similar magnitude and represented a similar increase over the stable 1000-2000 W/m²·K regime, *average* peak conductivity values at 9 W were approximately 10% higher than peak conductivity values at 8 W. This seems to at least indicate that the vaporization phenomena prior to dryout were similar. Most likely this also indicates that significant areas of the evaporator reached (or at least approached) the critical heat flux, and the largest fractions of working fluid were vaporized.

Although I was able to clearly demonstrate successful operation of the open-loop microevaporator setup, as well as experimentally verify the methods used to calculate effective thermal conductivity, the greatest disappointment of this study was my failure to develop quantitative performance trends or optimization schemes related to the evaporator *design*. Specifically, I did not see significant differences in performance based on columnation scheme, ceiling microtexture, or combination thereof. Neither did my design metrics of Columnar Area Percentage (CAP) and Column-to-Floor Ratio (CFR) show meaningful variation. This was primarily due to four factors:

- 1) Testing individual floor-ceiling combinations took much longer than anticipated. Specifically, I severely underestimated the time required to reach steady state at each power level. As a consequence, I was not able to *fully* test more than 15-20 out of the possible 240 floor-ceiling combinations.
- 2) For the vast majority of the floor-ceiling combinations tested, I spent far too much time developing full boiling curves (*i.e.* sweeping through all four performance regimes), rather than limiting my focus to the highest performance regimes (between Regime IV and dryout). Once I started with broad power level sweeps, I continued for the sake of consistency in my data and because I believed I would eventually observe significant differences in performance.
- 3) Again, due to time constraints, the vast majority of the tests were performed at only one mass flow rate. This limited the maximum heat flux I felt comfortable applying without risking thermal damage to the device clamps. As a consequence, I wasn't able to test enough devices under the high performance conditions for which they were designed.
- 4) When operating at or near full vaporization (Region IV), where devices would ultimately be designed to operate, extremely large fluctuations in effective conductivity create a kind of "noise" (though not truly noise in the technical signal-to-noise sense) which makes it very difficult to directly compare performance from one device to the next. One can certainly look at the mean and standard deviation, but meaningful comparative analysis is necessarily more challenging.

One major question that remains is the extremely transient nature of the 10,000-20,000 W/m²·K peaks. Because these peaks were always observed close to dryout, it is assumed that a large fraction of the working fluid was vaporized. Two theories are proposed to explain the short duration. The first (and most obvious) is related to the previous paragraph; if large areas of the evaporator were indeed at critical heat flux, it could be that these spikes simply preceded

widespread dryout. Consequently, thinning liquid films and particularly increased interline region caused a rapid (though unsustainable) increase in vaporization rate that, in turn, led to a dramatic rise in in-plane conductivity—precisely as intended by design and predicted by Fig. 5.9.

A second, less optimistic possibility is that, for whatever reason, more dryout was occurring near the data collection line on the nominally “cold” side of the evaporator (*i.e.* Line 2 in all the IR images). Though certainly not as likely, given the stochastic nature of vaporization, this is not impossible. If this were indeed the case, the temperature difference between Lines 1 and 2 would be “artificially” lowered, giving a false indication of a lower temperature gradient across the evaporator and thus a higher thermal conductivity. Had I used single data collection *points* (rather than lines) for calculating ΔT , this would be a greater likelihood. Indeed, this is the very reason for using lines (and average temperatures) in the first place. Still, given the irregularity of 10-20k peaks, this possibility cannot be ruled out.

One issue that should certainly be investigated is the effect, if any, of the cold sink temperature on the effective thermal conductivity. In nearly all of the quantitative tests performed, the cold sink of the original design (a notched metal cylinder embedded in an ice-water mixture) was not used. The reason for this was that, with the cold sink in place, the required flux just to achieve vaporization was so high that I feared melting the clamps during accidental dryout. Instead, a strip of paper towel was used to wick away the condensed liquid, but the condenser surface was cooled only by evaporation into the ambient air. Much like limiting the mass flow, this prevented me from testing the devices at higher fluxes that may have allowed complete vaporization and thus higher conductivities. In hindsight, perhaps the reason conductivities of 1000-2000 W/m·K were so stable while those in the 10,000-20,000 W/m·K range were so fleeting was because the internally-convected latent heat could not be removed from the *condenser* (which was not kept at fixed temperature), not because they could not be consistently removed from the *evaporator*.

8.7 Final Comments

By far the most significant achievement in this study was the successful demonstration of the microevaporator concept. Out-of-plane fluid evaporation and subsequent vapor convection was shown to enhance the effective thermal conductivity over solid conduction by one to two orders of magnitude. This corresponds to the range expected by numerical modeling.

A second noteworthy accomplishment was the development of a test platform and accompanying data analysis capable of obtaining accurate measurements of effective thermal conductivity for a microfluidic system involving phase change and internal convection. As discussed at the end of Chapter 7, this is no small feat, as precise experimental heat transfer measurements in real systems are notoriously difficult to obtain.

As discussed in Chapter 7, even easily interchangeable floor and ceiling device components took a very long time to test thoroughly. Ideally, I would have swept through a larger range of power levels at multiple flow rates for all approximately 240 floor and ceiling combinations. However, this was simply not possible due to time constraints. Nevertheless, the open-loop test setup proved extremely successful and will no doubt be useful to further work on this project. Moreover, the techniques used to estimate convective losses and calculate effective thermal conductivities were verified experimentally by analyzing the special case of solid

conduction over a wide range of heat fluxes. This will allow much easier data analysis for similar clamping systems in the future.

The columnated design of the vapor chamber provides a very efficient method of distributing liquid over the surface of the evaporator. I believe very strongly that this out-of-plane design is far superior to previous micro/meso scale loop heat pipes that merely relied on in-plane parallel channels. From a geometrical standpoint, regardless of the surface wetting characteristics, allowing liquid to spread radially from each column inherently increases the potential area of meniscus interline region. Although I was not able to determine statistically significant performance trends regarding columnation scheme or ceiling microtexture, I was able to demonstrate evaporator conductivities at or near the DARPA TGP goals with realistic mass flow rates in a device structurally and functionally analogous to the original μ C-LHP design.

Bibliography

Chapter 1

1. Lehmann, V., The Physics of Macopore Formation in Low Doped n-Type Silicon, *Journal of the Electrochemical Society*, vol. 140, no. 10, pp. 2836-2843, 1993.
2. Lehmann, V. and Goesele, U., Porous Silicon Formation: A Quantum Wire Effect, *Applied Physics Letters*, vol. 58, pp. 856-858, 1991.
3. Canham, L.T., Silicon Quantum Wire Fabrication by Electrochemical and Chemical Dissolution of Wafers, *Applied Physics Letters*, vol. 57, pp. 1046-1048, 1990.
4. Hoelke, A.D., Pilchowski, J., Henderson, H.T., Saleh, A., Kazmierczak, M., Gerner, F.M., and Baker, K., Coherent Macro Porous Silicon as a Wick Structure in an Integrated Microfluidic Two-Phase Cooling System, *SPIE Conference on Microfluidic Devices and Systems*, vol. 3515, pp. 154-162, 1998.
5. Hoelke, A.D., Henderson, H.T., Gerner, F.M., and Kazmierczak, M., Analysis of the Heat Transfer Capacity of a Micromachined Loop Heat Pipe, *Proceedings of the ASME Heat Transfer Division*, vol. 3, pp. 53-60, 1999.
6. Hamdan, M., Cytrynowicz, D., Medis, P., Shuja, A., Gerner, F.M., Henderson, H.T., Golliher, E., Mellott, K., and Moore, C., Loop Heat Pipe (LHP) Development by Utilizing Coherent Porous Silicon (CPS) Wicks, *Proceedings of the 8th ITherm Conference*, pp. 457-465, 2002.
7. Hamdan, M., Gerner, F.M., and Henderson, H.T., Steady State Model of a Loop Heat Pipe (LHP) with Coherent Porous Silicon (CPS) Wick in the Evaporator, *Proceedings of the 19th IEEE SEMI-THERM Symposium*, pp. 88-96, 2003.
8. Dhillon N.S., Pisano, A.P., Hogue, C.W., and Hopcroft, M.A., MLHP – A High Heat Flux Localized Cooling Technology for Electronic Substrates, *Proceedings of ASME IMECE*, pp. 621-630, 2008.
9. Park, K. and Lee, K., Flow and Heat Transfer Characteristics of the Evaporating Extended Meniscus in the Micro-Capillary Channel, *International Journal of Heat and Mass Transfer*, vol. 46, pp. 4587-4594, 2003.
10. Wang, H., Garimella, S.V., and Murthy, J.Y., Characteristics of an Evaporating Thin Film in A Microchannel, *International Journal of Heat and Mass Transfer*, vol. 50, pp. 3922-3942, 2007.

11. Dhavaleswarapu, H.K., Garimella, H.V., and Murthy, J.Y., Microscale Temperature Measurements near the Triple Line of an Evaporating Film, *Journal of heat Tranfer*, vol. 131, pp. 061501-1 – 06501-8.
12. Stephan, P. and Busse, C., Analysis of the Heat Transfer Coefficient of Grooved Heat Pipe Evaporator Walls, *International Journal of Heat and Mass Transfer*, vol. 35, pp. 383-391, 1992.
13. Pence, D.V., Reducing Pumping Power and Wall Temperature in Microchannel Heat Sinks with Fractal-Like Branching Channel Networks, *Microscale Thermophysics Engineering*, vol. 6, pp. 319-330, 2002.
14. Alharbi, A.Y., Pence, D.V., and Cullion, R.N., Fluid Flow Through Microscale Fractal-Like Branching Channel Networks, *Journal of Fluids Engineering*, vol. 125, pp. 1051-1057, 2003.
15. Wang, X.-Q., Majumdar, A.S., and Yap, C., Numerical Analysis of Blockage and Optimization of Heat Transfer Performance of Fractal-like Microchannel Nets, *Journal of Electronic Packaging*, vol. 128, pp. 38-45, 2006.
16. Wechsato, W., Lorente, S., and Bejan, A., Optimal Tree-Shaped Networks for Fluid Flow in a Disk-Shaped Body, *International Journal of Heat and Mass Transfer*, vol. 45, pp. 4911-4924, 2002.
17. Kandlikar, S.G., High Flux Heat Removal: A Roadmap of Challenges and Opportunities, *Heat Transfer Engineering*, vol. 26, pp. 5-14, 2005.
18. Carey, V.P. and Wemhoff, A., Thermodynamic Analysis of Near-Wall Effects on Phase Stability and Homogeneous Nucleation during Rapid Surface Heating, *International Journal of Heat and Mass Transfer*, vol. 48, pp. 5431-5445, 2005.
19. Carey, V.P., Padilla, J., and Gan, Y., Homogeneous Nucleation of Vapor at Preferred Sites During Rapid Transient Heating of Liquid in Micropassages, *Journal of Heat Transfer*, vol. 129, pp. 1333-1340, 2007.

Chapter 2

1. Carey, V.P., *Liquid-Vapor Phase-Change Phenomena*, Taylor and Francis Group, New York, 1992.

Chapter 3

1. Kandlikar, S.G., High Flux Heat Removal: A Roadmap of Challenges and Opportunities, *Heat Transfer Engineering*, vol. 26, pp. 5-14, 2005.

2. Kandlikar, S.G. and Balasubramanian, P., Experimental Study of Flow Patterns, Pressure Drop, and Flow Instabilities in Parallel Rectangular Minichannels, *Heat Transfer Engineering*, vol. 26, no. 3, 2005.
3. Mukherjee, A. and Kandlikar, S.G., Numerical Study of Growth of a Vapor Bubble during Flow Boiling of Water in Microchannel, *Microfluidics and Nanofluidics*, vol. 1, no. 2, 2005.
4. Peng, X.F., and Wang, B.F., Evaporation Space and Fictitious Boiling for Internal Evaporation of Liquid, *Science Foundation in China*, vol. 2, pp. 55-59, 1994.
5. Peng, X.F., Tian, Y., and Lee, D.J., Arguments on Microscale Boiling Dynamics, *Microscale Thermophysical Engineering*, vol. 6, no. 1, pp. 75-83, 2002.
6. Iida, Y., Okuyama, K., and Sakurai, K., Boiling Nucleation on a Very Small Film Heater Subjected to Extremely Rapid Heating, *International Journal of Heat and Mass Transfer*, vol. 37, no. 17, pp. 2771-2780, 1994.
7. Carey, V.P. and Wemhoff, A., Thermodynamic Analysis of Near-Wall Effects on Phase Stability and Homogeneous Nucleation during Rapid Surface Heating, *International Journal of Heat and Mass Transfer*, vol. 48, pp. 5431-5445, 2005.
8. Carey, V.P., Padilla, J., and Gan, Y., Homogeneous Nucleation of Vapor at Preferred Sites During Rapid Transient Heating of Liquid in Micropassages, *Journal of Heat Transfer*, vol. 129, pp. 1333-1340, 2007.
9. Webb, R.L. and Kim, N.H., *Principles of Enhanced Heat Transfer*, Taylor and Francis Group, New York, 2005.
10. Hummel, R.L. and Young, R.K., Improved Nucleate Boiling Heat Transfer, *Chemical Engineering Progress Symposium Series*, vol. 61, no. 59, pp. 264-267, 1964.
11. Xu, K. and Lloyd, J.R., Pool Boiling of FC-72: A Comparison of Two Thin Porous Coatings on Heat Transfer Enhancement, *ASME Heat Transfer Division*, vol. 376, part 2, pp. 279-288, 2005.
12. Carey, V.P., *Liquid-Vapor Phase-Change Phenomena*, Taylor and Francis Group, New York, 1992.
13. Hong, C-C, Murugesan, S., Kim, S., Beaucage, G., Choi, J-W, and Ahn, C.H., A Functional On-Chip Pressure Generator Using Solid Chemical Propellant for Disposable Lab-on-a-Chip, *Lab Chip*, vol. 3, pp. 281-286, 2003.
14. Lazar, M., Ambrovic, P., and Mikovic, J., Thermal Decomposition of Azo-bis-isobutyronitrile in the solid phase, *Journal of Thermal Analysis*, vol. 5, pp 415-425, 1973.

Chapter 4

1. Stroehle, D., On the Penetration of Gases and Water Vapour into Packages and Cavities and on Maximum Allowable Leak Rates, *15th Annual Reliability Physics Symposium*, pp. 101-106, 1977.
2. Tummala, R.R., Rymaszewski, E.J., and Klopfenstein, A.G. (eds.), *Microelectronics Packaging Handbook: Semiconductor Packaging Part II*, p. 888, 1997.
3. Chen, N., KSBonder: Anodic Bonding Characterization, *Marvell Nanolab Equipment Manual*, mod. 9.02, 2003.
4. Madou, M., *Fundamentals of Microfabrication*, CRC Press, Boca Raton, FL, 2002.
5. Jaeger, R.C., *Introduction to Microelectronic Fabrication*, Prentice Hall, Upper Saddle River, NJ, 2002.

Chapter 5

1. Cytrynowicz, D., Hamdan, M., Medis, P., Shuja, A., Henderson, H.T., Gerner, F.M., and Golliher, E., MEMS Loops Heat Pipe Based on Coherent Porous Silicon Technology, *Space technology and Applications International Forum (STAIF)*, pp. 220-232, 2002.
2. Suh, J., Shuja, A., Gerner, F.M., and Henderson, H.T., Nucleate Boiling inside the Evaporator of the Planar Loop Heat Pipe, *Proceedings of International Engineering Congress and Exposition (ASME IMECE)*, pp. 75-83, 2004.
3. Incropera, F.P. and DeWitt, D.P., *Fundamentals of Heat and Mass Transfer (Fifth Edition)*, John Wiley & Sons, New York, p. 129, 2002.
4. *Ibid.*, p. 8.
5. Dhavaleswarapu, H.K., Garimella, H.V., and Murthy, J.Y., Microscale Temperature Measurements near the Triple Line of an Evaporating Film, *Journal of heat Transfer*, vol. 131, pp. 061501-1 – 06501-8.
6. Park, K. and Lee, K., Flow and Heat Transfer Characteristics of the Evaporating Extended Meniscus in the Micro-Capillary Channel, *International Journal of Heat and Mass Transfer*, vol. 46, pp. 4587-4594, 2003.
7. Muwanga, R. and Hassan, I., Local Heat Transfer Measurements on a Curved Microsurface Using Liquid Crystal Thermography, *Journal of Thermophysics and Heat Transfer*, vol. 20, no. 4, pp. 884-894, 2006.
8. Höhmann, C. and Stephan, P., Microscale Temperature Measurements at an Evaporating Liquid Meniscus, *Experimental Thermal and Fluid Science*, vol. 26, pp. 157-162, 2002.

9. Buffone, C. and Sefiane, K., Temperature Measurement near the Triple Line during Phase Change Using Thermochromic Liquid Crystal Thermography, *Experiments in Fluids*, vol. 39, pp. 99-110, 2005.

Chapter 6

1. Madou, M., *Fundamentals of Microfabrication*, CRC Press, Boca Raton, FL, 2002.
2. Jaeger, R.C., *Introduction to Microelectronic Fabrication*, Prentice Hall, Upper Saddle River, NJ, 2002.
3. *Berkeley Microfabrication Laboratory Website*, Lab Manual: <http://microlab.berkeley.edu/text/labmanual.html>
4. *Marvell Nanofabrication Laboratory Website*, Lab Manual: <http://nanolab.berkeley.edu/labmanual/labmantoc.html>
5. Bean, K.E., Anisotropic Etching of Silicon, *IEEE Transactions on Electron Devices*, vol. 25, no. 10, pp. 1178-1185, 1978.
6. Williams, K.R. and Muller, R.S., Etch Rates for Micromachining Processes, *Journal of Microelectromechanical Systems*, vol. 5, no. 4, p. 257, 1996.

Chapter 7

1. Incropera, F.P. and DeWitt, D.P., *Fundamentals of Heat and Mass Transfer (Fifth Edition)*, John Wiley & Sons, New York, p. 538, 2002.
2. Lloyd, J.R. and Moran, W.R., Natural Convection Adjacent to Horizontal Surfaces of Various Planforms, *ASME Paper 74-WA/HT-66*, 1974.
3. Incropera, p. 551.
4. Churchill, S.W. and Chu, H.S.S., Correlating Equations for Laminar and Turbulent Free Convection from a Vertical Plate, *International Journal of Heat and Mass Transfer*, vol. 18, p. 1323, 1975.
5. Incropera, p. 557.

Appendix A: Detailed Discussion of Nucleate Boiling

Excerpted from author's M.S. Thesis:

A Self-Nucleating Evaporator: An Experimental Study of Low-Superheat Pool Boiling via Injection of Gaseous By-Products from the Thermal Decomposition of Azobis-Isobutyronitrile

University of California, Berkeley, 2008

(Numbered external references are found at the end of Appendix A)

II.2 Introduction to Phase Stability

As mentioned in passing earlier, there is a thermodynamic barrier associated with phase change. This barrier is not unlike the energy barriers associated with other chemical reactions. Consider the following classical example. Under certain conditions, a system of oxygen and methane will become more thermodynamically stable (*i.e.* have lower free energy) when combusted into carbon dioxide and water. However, although the reaction products are more stable than the reactants, activation energy in the form of heat or a spark is still necessary to initiate the reaction. Such activation energy allows the system to overcome the small barrier associated with the thermodynamically-favorable reaction.

The same is true of phase transformation. At temperatures below saturation, a liquid cannot undergo phase change because the resulting vapor would have a higher free energy. At temperatures above saturation, however, the opposite is true; phase change will lower the overall free energy and result in a more stable system. However, heating a liquid to its saturation temperature does not *necessitate* a phase change. Rather, the saturation temperature is merely the point at which phase change becomes thermodynamically *possible*. The *probability* of phase change is dependent on the activation energy required. The nature of this activation energy and the process by which it initiates phase change is the subject of the following sections.

II.2.1 P-v Diagrams

In order to address the notion of phase in a more graphical context, it is appropriate at this time to introduce the *P-v diagram*. The *P-v* diagram is a two-axis plot with specific volume (the reciprocal of density) on the *x*-axis and pressure on the *y*-axis (see Fig. 2.6). Since temperature is related to the product of pressure and volume (recall the ideal gas law), the *P-v* diagram provides a quick approximation of the phase of the system. For example, the left side of the plot corresponds to low temperatures and low specific volumes (*i.e.* high densities), so this is where we expect to find liquid states. Conversely, the right side corresponds to high temperatures and high specific volumes (*i.e.* low densities), so here we expect to find vapor states. Ubiquitous to the *P-v* diagram is the *saturation curve* or “vapor dome” as it is commonly called. Points along the left side of the curve correspond to saturated liquid states of the system, while points along the right side correspond to saturated vapor states. The intersection at the apex of the dome is the critical point. Thus points within the dome correspond to saturated systems consisting of both liquid and vapor phases.




Fig. 2.6 – Archetypical P - v diagram with saturation curve.

It is also instructive to plot isotherms on a P - v diagram. In elementary texts, the isotherms are often shown as horizontal within the saturation dome (see Fig. 2.7), indicating constant pressure (as well as temperature) during the transition from stable, single-phase liquid to stable, single-phase vapor. This also suggests that phase change (indicated by the abrupt change in slope as the isotherm enters the vapor dome) begins as soon as the saturation condition is reached.

If we instead plot isotherms predicted by the van der Waals equation of state and extend these curves into the vapor dome (see Fig. 2.8), even though it represents a nonequilibrium region, we can begin to address the issue of phase stability in a more mathematical sense.




Fig. 2.7 – P - v diagram with simplified (horizontal) isotherms within the vapor dome.




Fig. 2.8 – P - v diagram with van der Waal isotherms extending into saturation dome.

It can be shown using nonequilibrium thermodynamics that there are two criteria for phase stability:

$$c_v > 0 \quad (7)$$

$$\left(\frac{\partial P}{\partial v}\right)_T < 0 \quad (8)$$

Eq. 7 states that the molar specific heat at constant volume must be greater than zero (*i.e.* if you add thermal energy to the substance, its temperature increases). This is known as the *criterion of thermal stability*. Eq. 8 states that the partial derivative of pressure with respect to volume (at constant temperature) must be negative (*i.e.* if you increase the pressure on a substance, it shrinks). This is known as the *criterion of mechanical stability*. Since Eq. 7 is satisfied for virtually all substances, Eq. 8 is a necessary and sufficient condition for stability of phase.

At this point, we must acknowledge the somewhat transient nature of these “states”. Although we can speak of the pressure and density at a specific point on a P - v diagram, every location in a *real* system is at all times fluctuating around a nominal value of temperature, pressure, density, *etc.* For a state far from the saturation curve, these perturbations are much too small to affect the observed state. Take for example a stable, subcooled liquid. Deviations in molecular density are not likely to exceed the limits consistent with a liquid phase. However, for states near the saturation curve, perturbations in density may exceed these limits, resulting in localized regions where the liquid molecular density is lowered almost to that of saturated vapor. For this reason such fluctuations are often referred to as *heterophase fluctuations*. The consequences of these heterophase fluctuations will be discussed further in Section II.2.3.

II.2.2 Spinodal Definition and Metastable States

Let us now return to the P - v diagram in Fig. 2.8. Because we are plotting pressure as a function of volume, the partial derivative in Eq. 8 is nothing more than the slope of the isotherms. Thus anywhere the slope of the isotherm is positive represents an intrinsically unstable (and thus thermodynamically inaccessible) state. Notice also that each isotherm assumes exactly one minimum and one maximum within the vapor dome. Between these extrema the slope of the isotherms is positive, indicating an inaccessible state. This suggests constructing a locus of points consisting of all the extrema within the vapor dome (see Fig. 2.9). This curve is known as the *spinodal curve*. The portion consisting of all the minima (to the left of the critical point, in blue) is known as the *liquid spinodal*, while the portion consisting of all the maxima (to the right of the critical point, in red) is known as the *vapor spinodal*. These curves can be thought of as the intrinsic limits of thermodynamic stability. Construction of the spinodal curve also delineates two regions that lie outside the intrinsically unstable region but still within the nonequilibrium vapor dome. These are known as *metastable* regions (see Fig. 2.10). Points located between the saturated liquid curve and the liquid spinodal correspond to metastable “superheated” liquid states, while points located between the saturated vapor curve and the vapor spinodal correspond to metastable “supersaturated” vapor states.




Fig. 2.9 – Defining the spinodal curves using the criterion of mechanical stability (after [1]).




Fig. 2.10 – Using the spinodal curves to define metastable and unstable regions (after [1]).

II.2.3 Initiating Phase Change

Earlier we stated that the saturation temperature is the initial point at which phase change *may* occur. Likewise, we just showed that the spinodal temperature is the ultimate point at which phase change *must* occur. Thus we have established both lower and upper limits on the thermodynamics of phase change. Noting that the curves that define these limits form the boundaries of the metastable regions, it follows that phase change always occurs within the metastable regions. Experimentally, bubble formation within a superheated liquid is generally observed to occur over a range of temperatures within the metastable region. This seems quite natural based on the previous discussion on the effect of density fluctuations near the saturation curve. Indeed, these fluctuations are precisely the “activation energy” described earlier that is responsible for phase change. As a system is heated further and further into the metastable region, the difference in molecular density between superheated liquid and saturated vapor diminishes. Increasingly larger heterophase fluctuations will eventually exceed the density limits consistent with a liquid phase, resulting in the formation of small vapor embryos within the liquid.

II.3. Nucleation

The phenomenon associated with the formation of vapor embryos within a superheated liquid is commonly referred to as *nucleation*. Nucleation is further divided into two major types, homogeneous and heterogeneous. *Homogenous nucleation* refers to bubble formation *completely within* a superheated liquid (*i.e.* the bubble is surrounded on all sides by liquid). *Heterogeneous nucleation*, on the other hand, refers to bubble formation on a (usually solid) surface in contact with a superheated liquid. (Note: Although a completely analogous situation exists for the nucleation of liquid droplets in a supersaturated vapor undergoing cooling, I will primarily utilize language with reference to heating, as that is the primary focus of this study.) Because both the “driving force” and “activation energy” for each type of nucleation can vary considerably, the subsequent boiling processes have dramatic effects on the overall heat transfer.

II.3.1 Homogeneous Nucleation

As defined above, homogeneous nucleation is the formation of a vapor embryo within a surrounding superheated liquid. Once nucleation has occurred, the question then becomes: can the vapor embryo remain at equilibrium with the surrounding liquid? And if not, what are the conditions that determine whether the embryo will collapse or grow? We first consider the system shown in Fig. 2.11, with a spherical vapor embryo of radius r surrounded by superheated liquid at pressure P_l and temperature T_l . Note that because the liquid is assumed to be superheated, $P_l < P_{\text{sat}}(T_l)$. For equilibrium, the temperature of the vapor must be equal to T_l , but the curvature of the interface requires that the pressure inside the bubble be larger. Specifically, the Young-Laplace equation (for spherical geometry) requires

$$P_v = P_l + \frac{2\sigma}{r} \quad (9)$$




Fig. 2.11 – Homogeneous nucleation occurs when a vapor embryo forms completely surrounded by liquid (after [1]).

where σ is the interfacial tension. By requiring the chemical potential in the two phases to be equal and by integrating the Gibbs-Duhem equation for both the liquid and vapor phases, it can be shown [1, p. 138-140] that only embryos of radius r_e will be in equilibrium with the surrounding superheated liquid, where

$$r_e = \frac{2\sigma}{P_{\text{sat}}(T_l) \exp\{v_l[P_l - P_{\text{sat}}(T_l)]/RT_l\} - P_l} \quad (10)$$

Next, we define an appropriate availability function Ψ for the system and require (from basic thermodynamic considerations) that $d\Psi = 0$ for equilibrium and also that Ψ be a local minimum (*i.e.* concave up) for *stable* equilibrium. It can be shown [1, p. 141-144], however, that although $d\Psi(r_e) = 0$, Ψ is a negative quadratic function of r in the vicinity of r_e , making Ψ a local maximum at $r = r_e$. The loss of one molecule from the embryo will result in an embryo of radius $r < r_e$, a range where a further decrease in r will result in a further decrease in Ψ . Conversely, the addition of one molecule to the embryo will result in an embryo of radius $r > r_e$, a range where a further increase in r will again result in a further decrease in Ψ . Consequently, a bubble of radius r_e is in an *unstable* equilibrium with the surrounding liquid. This has a tremendous impact on the required conditions for phase change. As heterophase fluctuations in the superheated liquid produce tiny localized region of vapor, only those embryos with radius $r > r_e$ will spontaneously grow and result in homogenous nucleation of a vapor phase. Embryos of radius $r < r_e$ will spontaneously collapse and not result in a phase change.

The previous discussion indicates that the likelihood of homogeneous nucleation is largely dependent on the kinetics that govern the formation of vapor embryos within a metastable superheated liquid. In any superheated liquid, a vapor embryo consisting of n molecules will continuously gain molecules by evaporation and lose molecules by condensation at the liquid-vapor interface. The difference in the rates of these two processes will determine whether or not an embryo reaches the critical size to initiate homogeneous nucleation. In an attempt to quantify these rates, let us first assume an idealized distribution of various-sized embryos, whose relative concentration is a consequence of the availability associated with producing an embryo of that particular size. Thus we assume

$$N_n = N_l \exp \left[-\frac{\Psi(r)}{k_B T_l} \right] \quad (11)$$

where N_n is the number of embryos of n molecules at equilibrium per unit volume and N_l is the number of liquid molecules per unit volume. If we define $j_{n,e}$ and $j_{n,c}$ to be the number of molecules evaporating and condensing (respectively) per unit area at the interface of an embryo consisting of n molecules, then an equilibrium distribution can be established only if

$$N_n A_n j_{n,e} = N_{n+1} A_{n+1} j_{n+1,c} \quad (12)$$

In other words, the number of embryos which grow from size n to $n+1$ via evaporation is exactly balanced by the number of embryos which shrink from size $n+1$ to n via condensation. A real superheated liquid, however, is not expected to have an equilibrium distribution. If we now define J_n to be the difference between the sides of Eq. 12, we see that J_n represents the net flux of the number of embryos in “size space” from n to $n+1$. In other words, J_n represents the *excess* number of embryos of size n which grow to size $n+1$ via evaporation (which are not offset by embryos of size $n+1$ shrinking to size n via condensation). By invoking a few more assumptions (*i.e.* treating n as a continuous variable and asserting that $J(n) \equiv J$ is constant for a fixed temperature), it can be shown [1, p. 144-149] that

$$J = N_l \sqrt{\frac{3\sigma}{\pi m}} \exp \left\{ \frac{-16\pi\sigma^3}{3k_B T_l [\eta P_{\text{sat}}(T_l) - P_l]^2} \right\} \quad (13)$$

where

$$\eta = \exp \left\{ \frac{v_l [P_l - P_{\text{sat}}(T_l)]}{R T_l} \right\} \quad (14)$$

This J is the rate at which embryos grow from size n to size $n+1$ per unit volume of liquid. And because J is constant for all n (*i.e.* all embryo sizes), it is also true for embryos of critical size ($r = r_c$). This means that as J increases, the likelihood that a bubble will exceed critical size, grow spontaneously, and initiate homogeneous nucleation also increases. Although not particularly useful from an absolute “numerical” standpoint, the behavior of J provides tremendous insight into the overall phenomenon of homogeneous nucleation. In particular, note that most of the temperature-dependent quantities appear in the exponential term of Eq. 13. Thus even a slight change in temperature can have an enormous effect on the number of critically-sized embryos created. Typically, a change of 1°C can change J by three or four orders of magnitude. Thus, from a practical standpoint, although J is a monotonically increasing function of temperature, there will exist a narrow range of temperatures below which homogenous nucleation will essentially never occur, and above which it occurs immediately. The median temperature of this range is usually referred to as the *kinetic limit of superheat*. For water at atmospheric pressure, the kinetic limit is on the order of 300°C.

As the previous discussion highlights, the kinetic limit of superheat for “normal” conditions is usually much, much higher than the saturation temperature. Consequently, it is often said that *homogeneous nucleation is characterized by large values of superheat*. For boiling applications in which the temperature variation is to be minimized, it may be desirable to initiate phase change without approaching the kinetic limit required for homogeneous nucleation. This brings us to the other type of nucleation.

II.3.2 Heterogeneous Nucleation

In contrast to homogeneous nucleation, which occurs completely within a superheated liquid, heterogeneous nucleation occurs at the interface between a superheated liquid and a (usually solid) surface. This type of nucleation is quite common in many applications where heat is transferred from a hot solid object to a cooler surrounding liquid. In this case, there will be a temperature gradient in the liquid, with the hottest fluid located right at the solid-liquid interface. Once this liquid has been heated into the metastable region, it becomes thermodynamically possible for heterophase fluctuations to produce vapor embryos at the surface. In the idealized case of nucleation on a perfectly flat surface, the embryo takes the shape of a portion of a sphere (see Fig. 2.12), where the contact angle θ is dependent on the surface tensions of the system.




Fig. 2.12 – Heterogeneous nucleation occurs when a vapor embryo forms adjacent to a solid surface (after [1]).

At this point, the analysis governing the stability, growth, or collapse of the embryo mirrors that of the previous section, with a few minor modifications to account for geometrical differences. Following the same procedure, it can be shown [1, p. 169-173] that

$$J = \frac{N_l^{\frac{2}{3}}(1 + \cos \theta)}{2F} \sqrt{\frac{3F\sigma}{\pi m}} \exp \left\{ \frac{-16\pi\sigma^3}{3k_B T_l [\eta P_{\text{sat}}(T_l) - P_l]^2} \right\} \quad (15)$$

where η is identical to that given by Eq. 14 and F is a geometrical factor given by

$$F = \frac{2 + 3 \cos \theta - \cos^3 \theta}{4} \quad (16)$$

Since the procedure used to establish J is essentially the same for both homogeneous and heterogeneous nucleation, it is not surprising that the threshold superheat values are comparable. There is certainly some dependence on θ (with a slight shift to lower temperatures for $\theta > 70^\circ$ and to slight shift to higher temperatures for $\theta < 60^\circ$), but by and large, we should expect to see very little difference from the superheat values required for homogeneous nucleation. (Note: deviation from the homogeneous limit *will* become significant as $\theta \rightarrow 180^\circ$, but contact angles larger than 108° are almost never encountered in the real world [1].) The same can be said of nucleation occurring on protruding surfaces or in cavities of various shapes. Although there will be some geometrical dependence, the threshold temperature for heterogeneous nucleation will be comparable to that of homogeneous nucleation.

Although the analysis of the previous section is certainly legitimate, such high levels of superheat for heterogeneous nucleation are rarely observed experimentally. To boil a pot of water on the stove at atmospheric pressure, superheats of $10\text{-}15^\circ\text{C}$ are typically sufficient to initiate heterogeneous nucleation (once the bulk liquid has been heated to saturation). The source of disagreement can be found in our original idealization of a *vapor* embryo in contact with a solid surface. The presence of insoluble *gases* at the surface was not considered. In reality, most surfaces are covered with naturally-occurring or machine-formed scratches, pits, or other irregularities. Depending on their size and geometry (as well as the wettability of the solid), such cavities can trap insoluble gases when the surface is initially covered with liquid (see Fig 2.13). Because the advancing liquid front is convex (to maintain a consistent contact angle), the “nose” of the front will contact the opposite side of the cavity before the contact line reaches the bottom of the cavity (as long as $\theta > 2\beta$). This entrapped gas will allow heterogeneous vaporization to occur at the liquid-gas interface at much lower temperatures than would be required for “true” nucleation of a new vapor embryo. Furthermore, once the vaporization process is initiated, newly formed vapor can refill the cavities, allowing sustained low-superheat nucleation at the liquid-vapor interface. Consequently, it is often said that *heterogeneous nucleation is characterized by smaller values of superheat*.




Fig. 2.13 – a) The convex nose of an advancing liquid front will contact the opposite sidewall of the cavity if $\theta > 2\beta$. **b)** This will entrain insoluble gas to provide an initial embryo (after [1]).

As we did for the case of homogeneous nucleation, we must now determine the conditions for which a bubble will grow and emerge from the cavity. For the time being, we will idealize the cavity as conical with mouth radius R and cone angle 2β . Even with this idealization, the analysis is extremely complicated because, as the bubble grows, its radius of curvature does not vary monotonically and is, in fact, highly dependent on the relationship between the contact angle θ and the cone angle 2β . Nevertheless, if we assume that the superheat must exceed the equilibrium value *throughout* its growth, then we must require that the superheat exceed the equilibrium value of the *minimum* radius of curvature. Depending on the relationship between contact angle and cone angle, this minimum radius of curvature may occur at different points during bubble growth (*i.e.* at the “point” of the cavity, at the “mouth”, *etc.*). However, for most realistic cavity geometries and liquids (where $R/r \leq 1$ throughout growth), it can be shown [1, p. 180-183] that the required superheat for the site to be active is given by

$$\text{Superheat} \equiv T_l - T_{sat}(P_l) > \frac{2\sigma T_{sat}(P_l)v_{lv}}{h_{lv}R} \quad (17)$$

Noting again that the minimum radius of curvature experienced by a bubble as it fills the cavity is not necessarily equal to the cavity mouth radius R , we can say more precisely that a cavity will be active if the minimum radius of curvature (r_{\min}) imposed on the bubble during growth is greater than a critical value r^* given by

$$r^* = \frac{2\sigma T_{sat}(P_l)v_{lv}}{h_{lv}[T_l - T_{sat}(P_l)]} \quad (18)$$

This implies that as we increase the superheat, the value of r^* decreases and more sites will become active. This criterion is especially useful in heat transfer applications where cavity geometry is established by a particular surface roughness R . In this case, it is common to assume that the minimum curvature occurs at the effective mouth, in which case $r_{\min} \approx R$.

Having established the criteria for bubble growth up to and out of the cavity mouth, we now turn to bubble growth out of the cavity and into the surrounding superheated liquid. To establish nucleate boiling, not only must a bubble grow out of the cavity, it must grow large enough so that buoyancy and drag forces overcome surface tension forces and allow detachment from the solid surface. This distinction is necessary because the superheat in the surrounding liquid is not usually uniform. Instead, we typically observe a thin thermal boundary layer of superheated liquid (thickness δ_l) near the solid surface, with cooler (saturated or even subcooled) bulk liquid above. As bubbles grow and detach (a process known as *ebullition*), a downwash of colder liquid from above breaks up the thermal boundary layer (see Fig. 2.14a). This lowers the liquid temperature near the surface and prevents bubble growth beyond the cavity. Transient conduction from the solid into the liquid eventually thickens the thermal boundary layer, providing superheat for further vaporization from both the solid surface and the surrounding liquid (Fig. 2.14b). Eventually, the bubble grows large enough such that its buoyancy and drag overcome surface tension and release it from the surface. During release, a small amount of vapor (and/or insoluble gas) is left within the cavity, providing the initial embryo to start the ebullition process again (Fig. 2.14c). The escaping bubble creates a new downwash of colder liquid and the process repeats.




Fig. 2.14 – The ebullition process (after [1]).

The details of the ebullition cycle are complex, and it is still not exactly understood today. However, many models of varying complexity have been put forth to effectively predict the growth and release phenomena. Most notably Hsu postulated that an embryo will grow outside the cavity only when the liquid temperature equals or exceeds the required equilibrium superheat *over the entire surface* of the bubble. In other words, the bubble will not grow unless it is completely surrounded by a thermal boundary layer whose temperature exceeds the equilibrium superheat for the bubble at its current size. I will not repeat the mathematical analysis [1, p. 186-190], but instead state only the results. Namely, for a cavity of mouth radius r to be active, it must fall between r_{\max} and r_{\min} , where

$$\left\{ \begin{matrix} r_{\max} \\ r_{\min} \end{matrix} \right\} = \frac{\delta_t}{4} \left[1 - \frac{\theta_{sat}}{\theta_w} \left\{ \begin{matrix} + \\ - \end{matrix} \right\} \sqrt{\left(1 - \frac{\theta_{sat}}{\theta_w} \right)^2 - \frac{12.8\sigma T_{sat}(P_l)}{\rho_v h_{lv} \delta_t \theta_w}} \right] \quad (19)$$

and

$$\theta_{sat} \equiv T_{sat}(P_l) - T_{bulk} \quad (20)$$

$$\theta_w \equiv T_{wall} - T_{bulk} \quad (21)$$

This complicated expression is best understood by examining a qualitative plot (see Fig 2.15; note that the vertical scale is logarithmic). For sufficiently large values of wall superheat (*i.e.* to the right of the “nose”), there are two solutions for r corresponding to the maximum and minimum cavity mouth radii. Cavities whose mouth radii fall between these two values will be active. Not surprisingly, as the superheat increases, the square root term grows in magnitude, and thus the range of active cavity sizes also increases. On the other hand, as the superheat decreases, the terms inside the square root eventually vanish and r has only one solution. This corresponds to the *single* active cavity size that can sustain nucleate boiling at the *minimum* possible superheat. To the left of the nose, solutions are imaginary, indicating that such superheat levels cannot sustain nucleate boiling, regardless of cavity size. Notice also that the upper limit (r_{\max}) remains essentially constant with increasing superheat, while the lower limit (r_{\min}) decreases with increasing superheat. This indicates that increasingly smaller cavities

require increasingly larger superheat to activate. This result has tremendous implications at the microscale to be discussed in Chapter IV.




Fig. 2.15 – Graphical representation of Eq. 19 illustrating how active cavity size range varies with superheat. Note logarithmic scale on y-axis (after [1]).

II.3.3 Site Deactivation and Associated Phenomena

The preceding section discussed how heterogeneous nucleation sites are activated, as well as the conditions required to sustain steady ebullition (*i.e.* nucleate boiling) at those sites. Recall from our earlier discussion of boiling regimes that one of the primary characteristics (and indeed benefits) of nucleate boiling is the low superheat required for large heat fluxes. However, this characteristic is contingent upon the existence of entrapped vapor (and/or insoluble gases) within surface cavities and defects. This vapor/gas-liquid interface allows the vaporization process to proceed at low superheat because the high superheat associated with *new* embryo formation is no longer necessary.

There are, however, several ways in which nucleation sites can become deactivated. The simplest way occurs if the liquid *entirely* wets the surface to begin with. Referring back to Figure 2.13, if $\theta < 2\beta$ (*i.e.* for highly wetting fluids, extremely “liquid-philic” surfaces, and/or “obtuse” cavities), insoluble gases may not become entrapped within cavities at all. Highly pressurized systems may also result in deactivated nucleation sites. Increasing the pressure

increases the solubility of the ambient gas within the liquid and also drives liquid into cracks and cavities. In fact, during experiments where this was purposely done, the superheat required to initiate nucleation was found to be of the same order as that required for homogeneous nucleation [1, p. 178]. Of course, once nucleation was *initiated*, the required superheat to *sustain* nucleate boiling dropped considerably, since a vapor-liquid interface was then established.

Let us now return to more typical systems in which insoluble gases *are* initially entrapped within cavities prior to heating. It was stated earlier that during the bubble release stage of the (steady) ebullition cycle, a small amount of vapor is left within the cavity, which then serves as the embryo for the subsequent cycle. After several minutes or hours of boiling, however, the insoluble gases that initially provided the gas-liquid interface for heterogeneous nucleation may eventually be carried off completely, leaving only vapor within the cavities. If the surface is then cooled, this entrapped vapor will condense, allowing liquid to entirely wet the surface. Consequently, the superheat required for subsequent boiling cycles may be much higher than was required for the initial boiling cycle. This is sometimes referred to as *cyclic boiling hysteresis*. Because cavities with simple (*i.e.* non reentrant) geometry are especially susceptible to this phenomenon, it is common with most conventional surface-liquid combinations. Note, however, that the opposite effect can also occur. Because gas solubility generally decreases with temperature, gas that was initially dissolved in the liquid at low temperature may eventually come out of solution *within inactive cavities* once the liquid is heated toward saturation. This can, in turn, activate previously inactive nucleation sites.

Finally, let us examine the effects of subcooling. *Subcooled boiling* refers to boiling in which the temperature of the bulk liquid (*i.e.* liquid far outside the thermal boundary layer) is significantly below the saturation point. In this case, the downwash of substantially colder liquid during bubble release (see Fig. 2.14a) may cause the remaining vapor within the cavity to condense completely. Subcooling is often desirable from a heat transfer perspective because heat transfer coefficients for subcooled boiling are typically slightly higher than those for equivalent saturated boiling. However, one must be wary of site deactivation due to cold fluid downwash, which may then require high superheat for reactivation.

Chapter III. Nucleate Boiling: Relevant Past Work

III.1 Historical Pool Boiling Enhancement

Given the many desirable heat transfer characteristics of nucleate boiling, researchers have been attempting to enhance pool boiling since the early 1930s. Research has primarily progressed along four major avenues: abrasive treatments, non-wetting coatings, designed (three-dimensional) cavities, and porous surfaces. The historical summary presented in the following four sections draws heavily upon the first half of Chapter 11 of *Principles of Enhanced Heat Transfer* by Webb and Kim [2, Ch. 11]. For further detail on the research cited, please consult the original text.

III.1.1 Abrasive Treatments

Some of the earliest attempts to enhance pool boiling were focused on abrasive treatments. It was clear that surface defects and scratches played some part in the nucleation process, and researchers were quick to study those effects. In 1931, Jakob and Fritz first investigated the effects of surface finish on nucleate boiling performance. Most notably, these researchers found that surface roughening (using sand blasting or finely-machined crisscrossing grooves) led to a decrease in required superheat for the onset of nucleate boiling, but that such treatments only worked for a few hours. Between 1954 and 1962, many more studies were performed using flat surfaces roughened with emery paper of different coarseness or lapping compound. These studies confirmed the boiling coefficient increases and aging effects observed by Jakob and Fritz and also showed that intermittent subcooling dramatically reduced the effects. Together, these works demonstrated that boiling enhancement is strongly tied to the density of nucleation sites and motivated later researchers to investigate other means of creating artificial cavities.

In 1959, utilizing the newly-emerging technology of high-speed photography, Clark *et al.* confirmed that surface defects and scratches were indeed the location of active nucleation sites. By 1960, Griffith and Wallis had established a quantitative relationship between cavity mouth size and the superheat required for site activation. In a 1965 study on the effects of mechanical pitting on superheat, Hummel found only marginal improvement with pitting on a stainless steel surface. This suggested to him that intentional mechanical pitting did not establish more (or more stable) nucleation sites than a plain untreated stainless steel surface, perhaps pointing toward a microscale effect. In 1969, Chaudri continued the study of the aging effects of abrasive treatments and determined that, regardless of the abrading technique used, any benefits were completely mitigated after several hundred hours of operation. Whether this was due to degassing, fouling, or some other phenomenon was not immediately clear. However, the study confirmed the unavoidable aging effects in classically abraded surfaces, and consequently there was little sustained interest in that avenue thereafter.

III.1.2 Non-Wetting Coatings

The second major avenue of study in the enhancement of pool boiling has been the use of non-wetting surface coatings. In many respects, this was a logical progression based on the results of earlier studies concerning the gradual performance degradation observed in abraded surfaces. If abraded surfaces eventually lose their desirable nucleation characteristics due to degassing and complete cavity wetting, why not try to prevent such wetting?

In 1960, Griffith and Wallis were the first to suggest the use of non-wetting surface coatings to enhance pool boiling. Experimentally, they examined the boiling stability of a surface patterned with fine needle holes. They found that it was much easier to maintain nucleate boiling in such cavities if they were first coated with paraffin. A 1964 study by Hummel and Young [3] investigated the use of Teflon spots and Teflon-coated pits on stainless steel. Not only was there a significant decrease in the superheat required for nucleate boiling (approximately 5°F for the treated surface vs. 25°F for the untreated steel), there was also much less temperature variation *during* prolonged nucleate boiling (less than 0.2°F for the treated

surface vs. 2-3°F for the untreated steel). This suggests that non-wetting coatings not only help to activate more nucleation sites, but also tend to stabilize the ebullition cycle at those sites that have become active. Hummel further developed this technique in 1965 with a more consistent Teflon application process. By injecting a Teflon emulsion through an atomizer into a rising air stream, he was able to cover a stainless steel surface with Teflon spots approximately 0.25 mm in diameter at a density of 30-60 spots/cm². This surface decreased the superheat required for nucleation to less than 0.5K.

In 1967, Gaertner performed similar tests on a surface with artificial nucleation sites consisting of cavities coated with a low surface energy material. Like Hummel, he found that coated sites required lower superheat for activation and remained active longer than uncoated sites. In addition, he found that coating the entire surface (as opposed to just spots) does not enhance pool boiling. Instead, the bubbles tend to spread across the surface, coalesce, and eventually blanket large portions of the surface. In 1969, Vachone *et al.* also examined the effects of completely coating a metallic surface with Teflon. As opposed to Gaertner, however, they found that very thin coatings (8 μm was cited) improved performance, while thicker coatings (40 μm) tended to insulate the surface and worsen performance. Given the fabrication techniques used and the thinness of the 8 μm coating, it is the opinion of the present author that the improved performance may likely have been due to “patchiness” of the ultra-thin coating, creating an essentially spotted surface like that of Hummel.

A 1968 study by Bergles *et al.* explored the use of non-wetting surface coatings with fluids other than water. They found that highly-wetting fluids such as classical refrigerants were essentially unaffected by traditional surface coatings such as Teflon and that, consequently, non-wetting coatings are only useful for liquids with relatively large contact angles, such as water. In 1992, a tangentially-related but corroborating result was reported by Ko *et al.* in their examination of the effects of surface tension on boiling hysteresis. Recall from Section II.3.3 that hysteresis refers to the observed increase in superheat required for subsequent nucleation after a liquid has already been boiled and subcooled. Ko found that a decrease in surface tension (*i.e.* increased wetting) resulted in much larger hysteresis. This confirms Bergles’ conclusion that more wetting fluids tend to fully flood cavities and mitigate the nucleation promotion of non-wetting surface coatings.

III.1.3 Designed Cavities

The 1960 study by Griffith and Wallis not only demonstrated that superheat was related to cavity *mouth size*, it also showed that site stability was related to overall cavity *shape*. Based on theoretical grounds, it was understood that the liquid superheat required to maintain a stable vapor nucleus within a cavity was proportional to the reciprocal of the cavity radius. However, Griffith and Wallis took this a step further and considered cavity geometries that would allow for a *negative* vapor radius (see Fig. 3.1). This type of *doubly reentrant* cavity allows a stable vapor to exist within the cavity even in the presence of a *subcooled* liquid. Consequently, this type of geometry results in extremely stable nucleation sites. For further discussion of this phenomena, including the effect of liquid contact angle, see [1, p. 328-329]. Confirming this idea in 1961, Westwater constructed an array of reentrant cavities and observed that these cavities remained active long after “natural” (*i.e.* non-reentrant) cavities had flooded.




Fig. 3.1 – The geometry of a doubly reentrant cavity forces the bubble radius for most liquids to become negative when the interface lies within the cavity (after [2]).

Several other researchers throughout the 1960s and 70s found similar results for a variety of reentrant geometries. It is not uncommon for these types of enhanced surfaces to reduce superheat by a factor of ten or more, compared to a plain flat surface. For an excellent comparison of boiling curves for several well-known enhanced surfaces, see [1, p. 330]. The obvious shortcoming of designed (particularly reentrant) cavities is their complicated fabrication. This is generally done in one of two ways. In the first method, metal particles or molten droplets are fused together on a solid surface to produce a network of irregular cavities and passages. The second method involves first producing fin-like protrusions which are then partially flattened or bent and fused to the underlying surface to create fully reentrant cavities.

III.1.4 Porous Surfaces

As discussed above, the fabrication of regular fully reentrant cavities presents a great deal of added cost and complexity to most heat transfer applications, and as a result, recent studies have focused more upon porous surfaces for pool boiling enhancement. Some of the earliest work in this field was performed by Milton in the late 1960s and early 1970s. Patents issued in 1968, 1970, and 1971 described surfaces made of sintered metallic particles with void fraction ranging from 50-65%. Milton found that the optimum pore size is dependent on the fluid properties. In 1973, Gottzman *et al.* performed a series of tests on the Union Carbide material known as High-Flux. This sintered matrix of various sized particles (46% less than 44 μm in diameter, 54% between 44-74 μm in diameter) provided as much as 10x lower superheat for a prescribed heat flux. In addition, the material showed negligible aging effects and also increased the critical heat flux by 80% (for trichloroethylene).

Since then, numerous researchers and manufactures have explored many types of porous surfaces, usually fine metal powders sintered onto copper tubing. The most commonly used materials are copper, aluminum, and bronze, but other non-traditional materials have also been

studied. Typical particle sizes range between 12-500 μm in diameter, while the coating thickness is typically 0.25-2 mm. Another critical metric known as *porosity* (essentially the ratio of empty space to solid material) typically ranges between 40-70%. One additional complication with sintered powders is the required sintering temperature. Low-temperature sintering is highly desirable, so as not to anneal (and thus decrease the thermal conductivity of) the underlying copper tubing.

While much easier to manufacture and apply to various surfaces, porous materials do not provide the same site stability as fully reentrant cavities. For example, a 1982 study by Bergles and Chyu found full cavity flooding and significant hysteresis in High-Flux if the overlying liquid was sufficiently subcooled. In a more recent study in 2005, Xu and Lloyd [4] examined the effects of extended immersion of High-Flux *prior to heating*. Employing a new metric they called *Temperature Overshoot* ($TOS = T_{\text{boiling incipience}} - T_{\text{steady boiling}}$), they found that TOS increased with increasing immersion time (up to 24 hours). Thus even without cyclic boiling, many cavities became completely flooded by extended immersion alone. The degree of this hysteresis-like effect was highly dependent on liquid contact angle and was even more dramatic with subcooling. For an excellent summary of data and results for a variety of porous surfaces, please see [5].

III.2 Microscale Effects

The historical survey of enhanced nucleate boiling presented thus far has been restricted to *macroscale* pool boiling exclusively. In all the studies considered, the characteristic physical dimensions of the system (most often the evaporator surface length L) has been large compared to the bubble or capillary length scale L_b , defined as

$$L_b = \sqrt{\frac{\sigma}{g(\rho_l - \rho_v)}} \quad (22)$$

However, as discussed in the Introduction, the goal of this study is to stabilize the nucleate boiling regime in order to maximize heat transfer in *microscale* devices. Unfortunately, this has proven to be exceptionally difficult in practice. For one, the observable characteristics of phase change become dramatically different as hydraulic diameters decrease below approximately 100 μm . What's more, the fundamental physics of microscale nucleation and bubble growth are not entirely understood, and numerous divergent (and sometimes directly contradictory) theories have emerged in the literature to explain the phase change phenomena unique to the microscale. The next three sections summarize some of the experimental observations and relevant research in this field.

III.2.1 Single-Phase Cooling

The rapid development of microfluidic systems in the 1990s ushered in a dual wave of optimism in the fields of heat transfer and electrical engineering. The idea of fully-integrated refrigerant-filled microchannels intermeshed with CMOS seemed like the perfect long-term solution to the ever-growing thermal problem associated with denser, higher-powered integrated circuits and laser diodes. As a result, microfluidic systems for heat transfer quickly became a hot

topic, and researchers all over the world began reworking classical convection correlations to better model heat transfer in microchannels. Microchannels are naturally well-suited to cooling applications because of their large surface-to-volume ratio. However, they also create a host of new challenges centered upon the steady, efficient circulation of coolant to and from the heat-producing device. The following summary of early work in microchannel-based cooling draws heavily upon a 2005 “roadmap” by Kandlikar [6]. Please consult the original text for more detailed citation information.

In 1981, Tuckerman and Pease first proposed the concept of using microchannels for electronics cooling. Like all early efforts, they focused strictly on single-phase flow, and utilized the direct circulation of water in microchannels fabricated in silicon. Although their parallel microchannel heat sink was capable of dissipating 7.9 MW/m^2 with a substrate-to-(inlet)-liquid temperature difference of 71°C , the design resulted in very large pressure drops (200 kPa for plain microchannels, 380 kPa for pin-enhanced microchannels). In 1987, Phillips analyzed the heat transfer and fluid flow in microchannels and provided the first detailed equations for designing microchannel geometries. More recently in 2005, Kandlikar and Upadhye presented a set of equations for selecting optimal flow geometry for a prescribed pressure drop. Particularly, they examined the benefits of “split flows” (as opposed to unidirectional parallel channels, see Fig. 3.2) and developed optimization schemes for channel number and fin spacing. The decrease in pressure drop provided by the split flow concept led other researchers to seek nontraditional channel geometries that shorten the maximum flow length.




Fig. 3.2 – a) Traditional unidirectional parallel channels. **b)** “Split flow” channel geometry (after [6]).

One notable example presented by Colgan in 2005 utilized a single “microcavern” design containing a large array of staggered fins as opposed to a series of channels (see Fig. 3.3). Although the associated friction factors were higher than traditional channels, multiple inlet and outlet ports kept the maximum flow length under 3 mm and resulted in heat transfer coefficients as high as $130,000 \text{ W/m}^2$ with a pressure drop of less than 35 kPa.




Fig. 3.3 – Colgan’s (2005) novel “microcavern” heat sink design (from [6]).

III.2.2 Experimental Observations of Phase-Change Cooling

Although single-phase liquid cooling continues to draw interest from both academia and industry, the true “holy grail” of electronics cooling still focuses on two-phase boiling flow. There are several reasons for this, as enumerated by Kandlikar in his aforementioned “roadmap”:

1. The heat transfer coefficients in conventional flow boiling systems (> 3 mm hydraulic diameter) are very high compared to the corresponding single-phase values.
2. The required mass flow rates are reduced because of the use of latent heat in carrying the heat away, rather than just the sensible heat of the coolant (being limited by the available temperature rise in the coolant).
3. The heat removal process does not raise the temperature of the coolant as in the single-phase case, where the available temperature difference between the channel surface and the cooling fluid decreases along the flow length. This also results in a more isothermal evaporator surface, which may be critical for cooling highly temperature-sensitive electronics.

Unfortunately, the development of flow boiling microfluidic systems has been extremely limited by difficulties associated with 1) working fluid selection and 2) stable two-phase operation. With regard to the former, the low operating temperatures desired by electronics manufacturers preclude the use of water as a working fluid unless the system is operated at sub-atmospheric pressures. At the same time, experimental (as well as ecological and health safety) data on other refrigerants is not yet sufficiently developed to effectively steer the designer toward the proper working fluid.

With regard to the latter difficulty, namely maintaining stable two-phase operation, there are many recent studies describing complicated flow phenomena that are unique to the miniscale and microscale. Perhaps the most commonly observed phenomena have been the extreme pressure fluctuations and backflow associated with rapid bubble growth. A 2005 experimental study by Kandlikar and Balasubramanian [7] provides excellent visual confirmation of these

phenomena. Using high-speed photography (4,000-15,000 fps) of vaporization in parallel channels with a hydraulic diameter of 333 μm , they were able to observe nucleation along channel sidewalls, followed by extremely rapid bubble growth that quickly filled the entire channel width to form a large vapor slug. Interface velocities were as high as 3.5 m/s. Because of this rapid growth rate, the slugs do not “flow” along the channel (as in macroscale tubes), but instead fill the entire channel cross-section and expand along the channel in *both* directions (see Fig. 3.4 below). This naturally leads to backflow and enormous pressure drop fluctuations. Reverse slug growth even extended all the way into the inlet header in certain cases, causing severe flow maldistribution. A subsequent numerical simulation by Mukherjee and Kandlikar [8] confirmed these large interface velocities and backflow phenomena.




Fig. 3.4 – Violent vapor burst and backflow reported by Kandlikar and Balasubramanian (after [7]).

Jiang, *et al.* produced a series of papers in 1999 [9] and 2001 [10] that attempted to construct predictive boiling curves for two-phase flow in microchannels with significantly smaller hydraulic diameters (20-50 μm). Specifically, they were investigating the absence of (or at least non-traditional nature of) the macroscale “boiling plateau” that occurs during saturated nucleate boiling (see Fig. 3.5 below). Instead, they postulated three unique microscale boiling regimes, supported by slope changes in the boiling curve. In the first regime ($q/q_{CHF} < 0.4$), they observed single-phase liquid flow with occasional local nucleation, but these bubbles either collapsed or convected downstream without growth and did not appear to affect the single-phase boiling curve. In the second regime ($0.4 < q/q_{CHF} < 0.6$), bubbles larger than the channel’s width formed at the inlet header and only entered the channels when sufficient pressure was applied. This sporadic bubble injection produced a highly unstable flow, with random liquid droplets appearing in a periodic high-velocity vapor core. Finally, the third regime ($0.6 < q/q_{CHF} < 0.9$) produced a stable annular flow with active heat transfer at the liquid-vapor core interface. In

summary, they concluded that saturated nucleate boiling (and its associated boiling plateau) cannot occur in microchannels less than 100 μm in hydraulic diameter.




Fig. 3.5 – Boiling curves reported by Jiang *et al.* hypothesizing three distinct boiling regions based on changes in slope (from [10])

In 2001, Koo, *et al.* [11] presented a theoretical study of two-phase microchannel heat sinks with hydraulic diameters less than 100 μm . They, too, reported rapid phase change upon nucleation, resulting in an unstable vapor eruption that lacked traditional bubbly and slug flow regimes. Instead, they observed a vapor “onset point” that oscillated erratically in the longitudinal direction. A 2007 study by Lee *et al.* [12] reported similar instabilities in their own convective boiling experiments using methanol as the working fluid. Two unique flow patterns were observed. At lower heat fluxes and flow rates (0.86 W and 0.3 mL/hr, respectively), they observed an oscillating liquid/vapor interface (OI) that spanned the entire width of the channel (see Fig. 3.6). At higher fluxes and flow rates (4.8 W and 6.2 mL/hr), they observed spontaneous liquid-to-vapor burst (LB) that vaporized all of the liquid in the channel. At intermediate fluxes and flow rates (0.98 W and 0.6 mL/hr), they observed a combination of both. Utilizing in-channel MEMS temperature and pressure sensors, they were able to determine the unique frequencies of the OI and LB phenomena at each flux level and flow rate. In all cases, they found that the temperature and pressure oscillations were 180° out of phase.




Fig. 3.6 – Oscillating interface observed by Lee *et al.* (from [12]).

Another interesting observation has been the apparent location-dependence of homogeneous nucleation events, particularly in instances of extremely rapid transient heating. Recall from Section II.3.2 that location-dependence is typically a characteristic of *heterogeneous* nucleation (due to surface defects and cavities). Yet studies by O’Horo and Andrews [13] and Balss *et al.* [14] reported repeated *homogeneous* nucleation at specific locations on a heated surface. A particularly interesting result of a related study by Andrews and O’Horo [15] was the report of bubble formation by homogeneous nucleation near, but not at, the heater surface during rapid transient heating. As pointed out by Carey, this appears to contradict the expectation that homogeneous nucleation, if it occurs, should take place immediately adjacent to the surface itself, where the liquid is most highly superheated. In summary, these studies suggest that some mechanism either suppresses nucleation very near the wall or makes it more likely a certain distance away.

III.2.3 Competing Theories on Microscale Nucleation

The wide variety of phenomenological observations associated with phase change in microchannels has led a number of researchers to suggest theoretical explanations for the divergence from classical macroscale boiling models. Before introducing these theories, however, it is important to understand that universal consensus on this matter has by no means been established, and completely new theories (as well as modifications to existing theories) continue to populate the current literature.

Peng and Wang were among the first researchers to attribute entirely new phenomena to microscale boiling. In a 1994 study [16], they hypothesized the existence of a critical minimum “evaporating space” required for traditional nucleation. Below this critical space, no bubbles formed, despite the fact that the boiling curve indicated that they were indeed within the nucleate boiling regime. They further hypothesized the existence of what they called “microbubbles” not viewable by traditional means, and suggested a new dimensionless parameter

$$N_{mb} = \frac{h_{fg} a_v}{c \pi (v_l - v_v) q d_h} \quad (23)$$

where a_v is the thermal diffusivity of the vapor, v_l and v_v are evaluated at saturation conditions, d_h is the hydraulic diameter, and c is an empirical constant. If $N_{mb} \leq 1$, traditional nucleation is expected. If $N_{mb} > 1$, the result is a nonequilibrium state with highly efficient transport characteristics. Peng and Wang coined this nonequilibrium state “fictitious boiling”. In 1995, Cornwall and Kew suggested a different dimensionless number called the “confinement number”, which is simply a ratio of bubble length (see Eq. 22) to hydraulic diameter:

$$Co = \frac{1}{d_h} \sqrt{\frac{\sigma}{g(\rho_l - \rho_v)}} = \frac{L_b}{d_h} \quad (24)$$

Cornwall and Kew studied flows for $0.5 < Co < 10$ and reported four distinct flow regimes they termed isolated bubble (IB), confined bubble (CB), annular slug flow (ASF), and partial dryout (PD).

A 2002 study by Peng *et al.* [17] attempted to further quantify and support the “fictitious boiling” theory. Drawing upon Carey’s understanding of critical bubble size for growth and detachment, they again endorsed the concept that the channel size must be sufficiently large to allow bubbles to reach the critical radius r_c . If the channel is too small, the bubbles form and collapse, mimicking a boiling state without visible nucleation. In support of this concept, they made the following argument:

To maintain “traditional” boiling flow, the pressure impulse δP developed by a growing bubble must not “markedly alter” the overall system characteristics. Therefore $\delta P \ll \Delta P$, where ΔP is the total pressure drop of the flow. Assuming isothermal compressibility, *i.e.*

$$\beta = -\frac{1}{V} \left(\frac{\partial V}{\partial P} \right)_T \quad (25)$$

we can quickly derive the following:

$$\beta \delta P = -\frac{\delta V}{V} \sim -\frac{r^2}{r_h^2} \quad (26)$$

where r is the vapor embryo radius and r_h is the hydraulic radius ($d_h/2$) of the channel. Combining Eq. 26 with the pressure impulse requirement ($\delta P \ll \Delta P$), we get

$$r \ll r_h \sqrt{\beta \Delta P} \quad (27)$$

In summary, to maintain traditional boiling flow, the embryo radius must at all times be much less than the quantity on the right side of the inequality above. Substituting typical, but conservative values into the right side of the inequality (water at atmospheric pressure, $\Delta P = 1$ bar, $r_h = 1$ mm), we find that $r \ll 1$ μm . This presents a serious problem, however, because Carey’s formulation for similar conditions requires $r_c \sim 1$ μm . Thus this pressure-based

argument further supports the idea that a minimum evaporation space is required for traditional nucleation.

In the same study, Peng *et al.* discussed another possible mechanism for non-traditional boiling that has frequently been termed “cluster theory”. In essence, this theory postulates the existence of clusters of high-energy nonequilibrium molecules that behave as large Brownian particles. These clusters can coalesce and grow to become true nucleated bubbles, or they can impact channel walls or other low energy clusters and dissociate. Based on the results of Brownian theory (see original text for derivation), the following relation is established:

$$\overline{(x - x_0)^2} = \frac{2k_B T}{\alpha} t \quad ; \quad \alpha = 6\pi r \eta \quad (28)$$

Here, the left side is a representative “traveling distance” for the Brownian particle, t is a representative “traveling time”, and α is a damping coefficient related to the Brownian particle radius r and the dynamic viscosity η based on Stokes’ Law. Again, substituting typical, but conservative values for superheated water ($T = 120^\circ\text{C}$, $r = 10^{-8}$ m) and utilizing an experiment-based observation that nucleation typically occurs around 30 sec after the necessary superheat is achieved, they found that the characteristic traveling distance is 0.0672 mm. This implies that an energized cluster must coalesce with other energized clusters within 30 sec and within a distance of 67 μm in order to form a nucleus. Collisions with channel walls or other non-energized clusters within this distance will tend to dissociate the cluster and prevent nucleation. While this “cluster theory” is certainly highly speculative, what *is* interesting is that the characteristic traveling distance it predicts (67 μm) is extraordinarily close to the transition length scale between macroscale and microscale boiling phenomena.

In 1994, Iida *et al.* [18] confirmed homogeneous nucleation conditions on thin platinum films subjected to heating rates of up to 9.3×10^7 K/s. They found that, although the wall temperature at boiling incipience increases with increased heating rate, the wall temperature eventually saturates (levels off) at the highest rates and that these temperatures correspond to the expected homogeneous nucleation temperatures. They did not, however, attribute this to traditional homogeneous nucleation. Recall from Section II.3.1 that subsequent bubble growth after homogeneous nucleation is normally attributed to heat transfer *from* the liquid-vapor interface *to* the surrounding fluid. Based on what they perceived to be “spontaneous” (not homogeneous) nucleation, they instead concluded that vaporization by excess energy in the superheated liquid *prior to nucleation* is the dominant mechanism for bubble growth. This is very close to the central concept of the next theory to be presented, which the current author believes to be the most persuasive.

In response to the report by Andrews and O’Horo of homogeneous nucleation occurring near, *but not at*, the heater surface, many researchers sought to understand the mechanism of nucleation suppression in the extreme near-wall region. A likely candidate appeared to be force interactions between molecules of the liquid and molecules of the solid affecting the intrinsic thermodynamic stability limits of the fluid. A 2005 study by Carey and Wemhoff [19] indicates that the wall-fluid attractive forces result in a rise in the pressure very near the solid surface (see Fig. 3.7). Even more significantly, their model predicts that the local spinodal temperature (see

Section II.2.2) increases rapidly in the region within a few nanometers of the wall (see Fig. 3.8). As a result, the spinodal condition (*i.e.* kinetic limit of superheat) will be first achieved not at the wall itself, where the fluid is hottest, but rather at a location a few nanometers away from the surface, where the spinodal temperature is lower.




Fig. 3.7 – Near-wall effects on pressure reported by Carey and Wemhoff (after [20]).




Fig. 3.8 – Near-wall effects on spinodal temperature reported by Carey and Wemhoff (after [20]).

This is perhaps best understood in the case of rapid transient heating (see Fig. 3.9), the subject of a 2007 study by Carey *et al.* [20]. As heat is conducted from the hot surface into the liquid, the transient linear temperature profile pushes further and further into the liquid until it intersects the spinodal temperature curve at a location z^* some distance away from the wall. Homogeneous nucleation will then first occur at this location, not at the wall, where the spinodal temperature has been artificially elevated by near-wall effects.




Fig. 3.9 – Nucleation initiated a small distance away from heated wall due to elevated spinodal temperature in the near-wall region (after [20]).

For most “real-world” surfaces (with characteristic roughness on the scale of microns to tens of microns), such nanoscale near-wall effects would be difficult to separate from microscale surface roughness effects. Consequently, homogeneous nucleation conditions are not typically observed when boiling a pot of water on the stove. On the other hand, however, the precision microfabrication of microfluidic systems often results in surfaces with submicron or even nanoscale characteristic roughness; in this case, such dimensions are comparable to that of the near-wall affected region. In such circumstances, propose Carey *et al.*, the near-wall region would presumably follow the contours of the surface (see Fig. 3.10), allowing a protrusion of lowered T_{spin} “bulk” liquid to extend down into a relatively “deep” cavity. Because this conical protrusion is nearly surrounded by solid surface, preferential (*i.e.* multi-direction) conduction to this low T_{spin} fluid during transient heating would cause it to exceed its low bulk spinodal temperature before fluid at other “generic” near-wall locations reach their elevated spinodal temperatures (*e.g.* point A will nucleate before points B and C). Numerical simulations of transient heating in idealized conical cavities further supported this theory. As summarized by Carey *et al.*, “This suggests that a small cavity containing a protrusion of low T_{spin} fluid would be a preferred homogeneous nucleation site in rapid transient heating.”

This theory provides a very persuasive explanation for the seemingly puzzling location-dependence of homogeneous nucleation observed by O’Horo and Andrews and Balss. In addition, it also predicts the existence of highly-superheated liquid in the near-wall regions immediately surrounding the homogeneously nucleated vapor embryo, *prior to nucleation*. Because nucleation was suppressed in these regions by elevated spinodal temperatures, there is tremendous excess thermal energy available to promote explosive bubble growth once nucleation of the low- T_{spin} protrusion has occurred.




Fig. 3.10 – A protrusion of low- T_{spin} fluid allows preferential conduction and explains location-dependence of spinodally-limited (homogeneous) nucleation (after [20]).

III.2.4 Attempts to Stabilize Nucleate Boiling

Having presented several competing theories to explain the various boiling phenomena unique to the microscale, it is now appropriate to discuss some of the approaches researchers have used in an attempt to curb the numerous undesirable effects of microscale vaporization, most notably large pressure fluctuations and violent vapor bursts. These efforts have primarily focused upon two central concepts: preventing fluid backflow and promoting heterogeneous nucleation using artificial nucleation sites.

As discussed in Section II.2.2, full-channel-width vapor slug growth in both directions (due to extremely rapid vaporization) has caused serious problems with backflow in many parallel channel flows. Kandlikar and Balasubramanian [21] first attempted to counter this trend by introducing pressure restrictors at the channel inlet. These pressure restrictors were essentially just microscale orifice plates that decreased the cross-sectional area of the channel inlet to raise the pressure necessary to generate backflow. In addition, they drilled a series of small holes (5-20 μm in diameter) along the channel length to serve as artificial nucleation sites. While neither approach worked by itself, when used in conjunction, backflow was completely eliminated. In addition, the artificial nucleation sites generated vapor bubbles that did not fill the entire channel cross-section, but instead detached and flowed along the channel in the intended direction.

A 2002 study by Zhang *et al.* [22] utilized ion-implanted in-channel thermistors to compare boiling conditions in both plain-walled and “enhanced” microchannels containing both notched sidewalls (two-dimensional “cavities”) and etched cavities in the channel “floor”. Although they did observe significantly lowered superheat for transition to two-phase flow in enhanced channels with 44 μm hydraulic diameter, a boiling plateau was not seen because annular flow quickly transitioned to mist flow after boiling incipience. They also investigated the effects of surface tension by varying the concentration of Triton X-100 surfactant from 200-10,000 ppm but found little to no effect on the required superheat. It should be stated, however, that in all cases, the sidewall notches failed to become active nucleation sites. This is, perhaps, an indication that the hydrophilic native oxide increased surface wetting and eliminated any entrapped gases to begin with, regardless of the concentration of the surfactant.

In 2005, Kuo *et al.* [23] improved on the notched sidewalls of Zhang, *et al.* by utilizing DRIE microfabrication to construct enhanced microchannels with two-dimensional reentrant cavities along the sidewalls (see Fig. 3.11 for SEMs and dimensioned schematic). Confirming the 1960 proposal by Griffith and Wallis, they reported that mouth diameter controls the required superheat for the onset of nucleation, while cavity shape determines the stability once boiling starts. More significantly, however, was their observation of both bubbly and slug flow, in addition to the usual single phase and annular flow. Although their channels fall toward the larger end of the microchannel spectrum (200 μm wide), stable bubbly flow has been quite difficult to achieve in any case. In addition, they found that the number of active nucleation sites increases with increased mass flow rate but shifts the nucleate boiling zone toward the entry region.




Fig. 3.11 – SEM micrograph of DRIE reentrant sidewall cavities by Kuo *et al.*, along with relevant dimensions (after [23]).

References for Appendix A

1. Carey, V.P., *Liquid-Vapor Phase-Change Phenomena*, Taylor and Francis Group, New York, 1992.
2. Webb, R.L. and Kim, N.H., *Principles of Enhanced Heat Transfer*, Taylor and Francis Group, New York, 2005.
3. Hummel, R.L. and Young, R.K., Improved Nucleate Boiling Heat Transfer, *Chemical Engineering Progress Symposium Series*, vol. 61, no. 59, pp. 264-267, 1964.
4. Xu, K. and Lloyd, J.R., Pool Boiling of FC-72: A Comparison of Two Thin Porous Coatings on Heat Transfer Enhancement, *ASME Heat Transfer Division*, vol. 376, part 2, pp. 279-288, 2005
5. Webb, R.L. and Kim, N-H, *Principles of Enhanced Heat Transfer*, Taylor and Francis Group, New York, 2005, pp. 403-405.

6. Kandlikar, S.G., High Flux Heat Removal – A Roadmap of Challenges and Opportunities, *Heat Transfer Engineering*, vol. 26, no. 8, pp. 5-14, 2005.
7. Kandlikar, S.G. and Balasubramanian, P., Experimental Study of Flow Patterns, Pressure Drop, and Flow Instabilities in Parallel Rectangular Minichannels, *Heat Transfer Engineering*, vol. 26, no. 3, 2005.
8. Mukherjee, A. and Kandlikar, S.G., Numerical Study of Growth of a Vapor Bubble during Flow Boiling of Water in Microchannel, *Microfluidics and Nanofluidics*, vol. 1, no. 2, 2005.
9. Jiang, L., Wong, M., and Zohar, Y., Phase Change in Microchannel Heat Sinks with Integrated Temperature Sensors, *Journal of MEMS*, vol. 8, no. 4, pp. 358-365, 1999.
10. Jiang, L., Wong, M., and Zohar, Y., Forced Convection Boiling in a Microchannel Heat Sink, *Journal of MEMS*, vol. 10, no. 1, pp. 80-87, 2001.
11. Koo, J-M, Jiang, L., Zhang, L., Zhou, P., Banerjee, S., Kenny, T., Santiago, J., and Goodson, K., Modeling of Two-Phase Microchannel Heat Sinks for VLSI Chips, *MEMS 2001*, pp. 442-426, 2001.
12. Lee, M., Lee, Y-K, and Zohar, Y., Two-Phase Flow Oscillations in Microchannel Convective Boiling, *MEMS 2007*, pp. 635-638, 2007.
13. O'Horo, M.P. and Andrews, J.R., Initial Stages of Vapor Bubble Nucleation in Thermal Ink Jet Processes, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, J. Bares, Editor, vol. 2413, pp. 182-188, 1995.
14. Balss, K.M., Avedisian, C.T., Cavicchi, R.E., and Tarlov, M.J., Nonosecond Imaging of Microboiling Behavior on Pulsed-Heated Au Films Modified with Hydrophilic and Hydrophobic Self-Assembled Monolayers, *Langmuir*, vol. 21, no. 23, pp. 10459-10467, 2005.
15. Andrews, J.R. and O'Horo, M.P., High-Speed Stroboscopic Systems for Visualization of Thermal Inkjet Processes, *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, J. Bares, Editor, vol. 2413, pp. 176-181, 1995.
16. Peng, X.F., and Wang, B.F., Evaporation Space and Fictitious Boiling for Internal Evaporation of Liquid, *Science Foundation in China*, vol. 2, pp. 55-59, 1994.

17. Peng, X.F., Tian, Y., and Lee, D.J., Arguments on Microscale Boiling Dynamics, *Microscale Thermophysical Engineering*, vol. 6, no. 1, pp. 75-83, 2002.
18. Iida, Y., Okuyama, K., and Sakurai, K., Boiling Nucleation on a Very Small Film Heater Subjected to Extremely Rapid Heating, *International Journal of Heat and Mass Transfer*, vol. 37, no. 17, pp. 2771-2780, 1994.
19. Carey, V.P. and Wemhoff, A., Thermodynamic Analysis of Near-Wall Effects on Phase Stability and Homogeneous Nucleation during Rapid Surface Heating, *International Journal of Heat and Mass Transfer*, vol. 48, no. 25-26, pp. 5431-5445, 2005.
20. Carey, V.P., Padilla, J., and Gan, Y., Homogeneous Nucleation of Vapor at Preferred Sites During Rapid Transient Heating of Liquid in Micropassages, *Journal of Heat Transfer*, vol. 129, no. 10, pp. 1333-1340, 2007.
21. Kandlikar, S.G. and Balasubramanian, P., Experimental Study of Flow Patterns, Pressure Drop, and Flow Instabilities in Parallel Rectangular Minichannels, *Heat Transfer Engineering*, vol. 26, no. 3, 2005.
22. Zhang, L., Wang, E.N., Koo, J-M, Jiang, L., Goodson, K.E., Santiago, J.G., and Kenny, T.W., Enhanced Nucleate Boiling in Microchannels, *MEMS 2002*, pp. 89-92, 2002.
23. Kuo, C., Koşar, A., Jensen, M.K., and Peles, Y., Boiling in Enhanced Surface Microchannels, *ASME Heat Transfer Division*, vol. 376, part 2, pp. 377-384, 2005.

Appendix B: Self-Nucleating Evaporator (SNE) Project

Excerpted from author's M.S. Thesis:

A Self-Nucleating Evaporator: An Experimental Study of Low-Superheat Pool Boiling via Injection of Gaseous By-Products from the Thermal Decomposition of Azobis-Isobutyronitrile

University of California, Berkeley, 2008

(For internal references from Chapters I-III, please see Appendix A. Numbered external references can be found at the end of Appendix A.)

Chapter IV. Self-Nucleation Concept

IV.1 Basic Principle

As discussed in the previous chapter, high-superheat site reactivation (and its associated vapor burst) has thus far hindered or prevented entirely the utilization of the nucleate boiling regime in microscale heat transfer systems. Moreover, gradual degradation of various types of enhanced surfaces has also limited long-term performance in many macroscale systems. The goal of this study, then, is to show how a self-nucleating evaporator surface can passively reactivate flooded sites at low superheat, thus stabilizing nucleate boiling and maximizing heat transfer efficiency.

The guiding principle of this concept is the dramatic role played by entrapped gases during heterogeneous nucleation (see Section II.3.2). Recall that insoluble gases entrapped within surface cavities allow heterogeneous nucleation to occur at much lower superheat than homogeneous nucleation because a gas/liquid interface is already present within the cavity. However, if a cavity becomes completely “flooded” for any reason (see Section II.3.3), site reactivation often requires the large superheat associated with homogeneous nucleation. To prevent this excessive superheat, a self-nucleating surface *injects* insoluble gases into deactivated nucleation sites, allowing low-superheat heterogeneous nucleation to continue. Such coerced reactivation has benefits for many heat transfer applications. Not only does it minimize the maximum wall temperature and create a more uniform temperature distribution across the evaporator surface, it also helps to prevent vapor burst, transition to film boiling, and surface dryout.

IV.2 Azobis-isobutyronitrile (AIBN)

The key to any self-nucleating surface is the injection of insoluble gases into deactivated nucleation sites. While simple in theory, such a system requires not only an electro-chemical reaction of some kind, but also some method for determining *which* sites have become deactivated. In addition, the system should be passive, self-regulating, and possess a functional lifetime appropriate for its application. While a few alternative systems will be discussed briefly in Section IV.4, the majority of this study will be devoted to the use of azobis-isobutyronitrile, or AIBN as it's commonly called.

IV.2.1 Physical and Chemical Properties

Azobis-isobutyronitrile (AIBN) is a chemical compound often used as an initiator for a variety of free radical polymerizations and also as a foamer in plastics and rubber. Its official IUPAC name is 2,2'-azobis(2-methylpropionitrile), with molecular formula $C_8H_{12}N_4$ and structure shown in Fig. 4.1. Its molar mass is 164.21 kg/kmol and its melting point is between 100-105°C. Physically, AIBN is a white, odorless crystalline solid, usually purchased as a fine white powder. Although considered a flammable solid, it is often used as an initiator instead of benzoyl peroxide because the risk of explosion is smaller. AIBN is soluble in both methanol and ethanol but has extremely low solubility in water.




Fig. 4.1 – Molecular structure of azobis-isobutyronitrile (AIBN).

IV.2.2 Thermal Decomposition

The most common chemical reaction associated with AIBN is its thermal decomposition. Like other azonitriles, the AIBN molecule is symmetric about the azo bridge (the double-bonded nitrogens), the most labile functional group. The azo bridge is easily thermally cleaved (at approximately 64°C), liberating gaseous nitrogen and leaving behind stabilized free radicals (see Fig. 4.2a). This reaction is the basis of the commercial utility of azonitriles as a source of free radical initiators and, to a lesser extent, as a source of nitrogen gas in foam blowing. According to the chemical equation, the AIBN will be reduced to 87% of its original weight once all of the nitrogen has been released. An experimental study by Hong et al [24] showed that the temperature was approximately 100°C when the remaining weight reached 87% (see Fig. 4.2b), indicating that that nitrogen effluence occurred over the range of approximately 64-100°C. Upon further heating, a second thermal decomposition of the remaining free radicals occurs near the melting point.

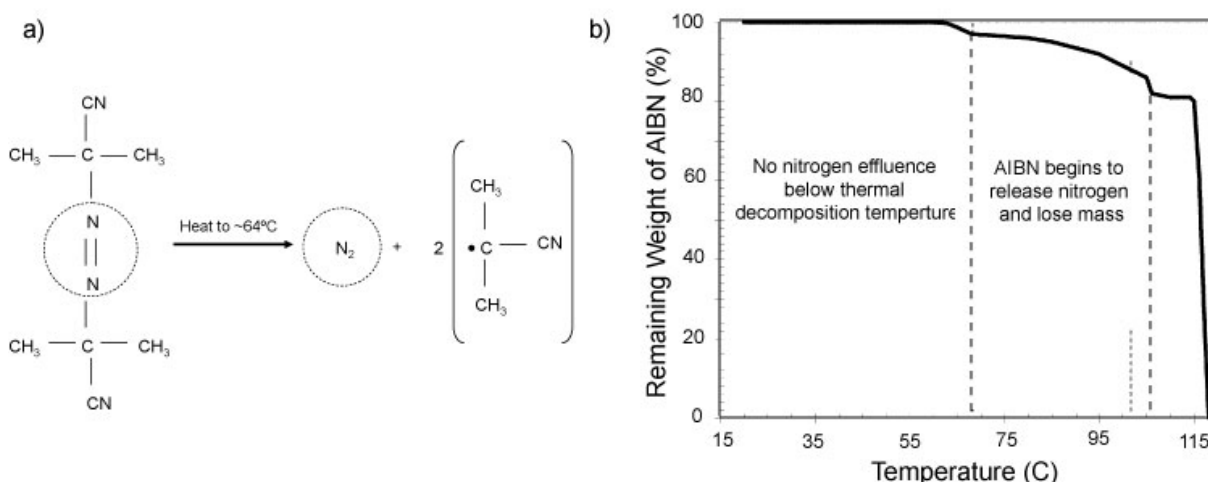


Fig. 4.2 – a) Thermal decomposition reaction of AIBN. **b)** Temperature range of decomposition (from [24]).

IV.2.3 Matrix Encapsulation

Because AIBN by itself is a fine white powder, some kind of polymer encapsulation is required to create a practical surface for evaporation. This polymer matrix is also necessary to confine the solid free radicals that remain upon thermal decomposition. Several criteria must be met when selecting an appropriate polymer. First, the polymer selected must be compatible with the working fluid (*i.e.* non-reactive, insoluble, *etc.*) and must adhere to both AIBN and the desired substrate. Second, it must cure at a temperature well below the first thermal decomposition temperature of AIBN but must also remain stable (*i.e.* not melt) at temperatures required for said decomposition. Finally, the polymer must have relatively low mechanical strength (or sufficient porosity) to allow liberated nitrogen molecules to break through the matrix and diffuse to the solid/liquid interface of deactivated cavities.

Fortunately, many polymers fulfill these requirements, and several, in fact, (such as spin-on Teflon, and SU-8 photoresist) are materials commonly employed in microfabrication. Most notably, AIBN has been successfully used as a thermally-activated solid propellant for microfluidic devices [24]. In this study, an underlying thin-film gold heater was used to heat the Teflon-encapsulated AIBN, thereby releasing nitrogen into microchannels to pump various liquids around the chip. Even using a crude, poorly-described “screen-printing” process, these researchers were able to successfully pattern the AIBN to a resolution of 1 mm x 1 mm with a thickness of 75 μm . With further study and better fabrication techniques, it would appear that “fully” microscale structures (*e.g.* AIBN-coated cavities *within* microchannels) are achievable.

IV.3 Site Reactivation – Two Narratives

At this point, I would like to offer two different narratives to describe the reactivation process of a deactivated nucleation site. The first narrative describes what happens at a “normal” site (*i.e.* one *not* treated with a coating of AIBN), while the second describes what happens at a site that *has* been treated with an AIBN-polymer coating. For the purpose of familiarizing the

reader with fluid and surface temperatures relevant to this study, I will describe a reactivation scenario that mirrors the conditions of the experimental sections (see Chapter V). For both narratives, we will assume a starting condition of stable nucleate boiling of a refrigerant whose saturation temperature is 34°C at atmospheric pressure. For stable nucleate boiling, a superheat of 10-15°C is typically expected (see Section II.3.2), so we will assume a uniform wall temperature of 50°C (see Fig. 4.3a). At this point, we then assume that a particular nucleation site becomes completely wetted and thus deactivated.

IV.3.1 Site Reactivation without Self-Nucleation

Deactivation of the nucleation site eliminates one mode of heat transfer and thus causes an immediate decrease in the local heat transfer coefficient. This, in turn, leads to a temperature rise at the surface and the formation of a localized hot spot beneath the deactivated site (Fig. 4.3b). This hot spot will grow and intensify until it develops sufficient superheat near the surface to initiate homogeneous embryo formation via preferential conduction. Local wall temperatures could easily reach 100°C, which would (at best) create large temperature gradients across the evaporator surface and could (at worst) damage or destroy underlying electronics. What's more, once an embryo has been nucleated (Fig. 4.3c), there is so much excess superheat in the surrounding liquid that vaporization at the liquid-vapor interface will proceed *extremely* quickly, resulting in violent vapor burst or complete surface dryout (Fig. 4.3d).

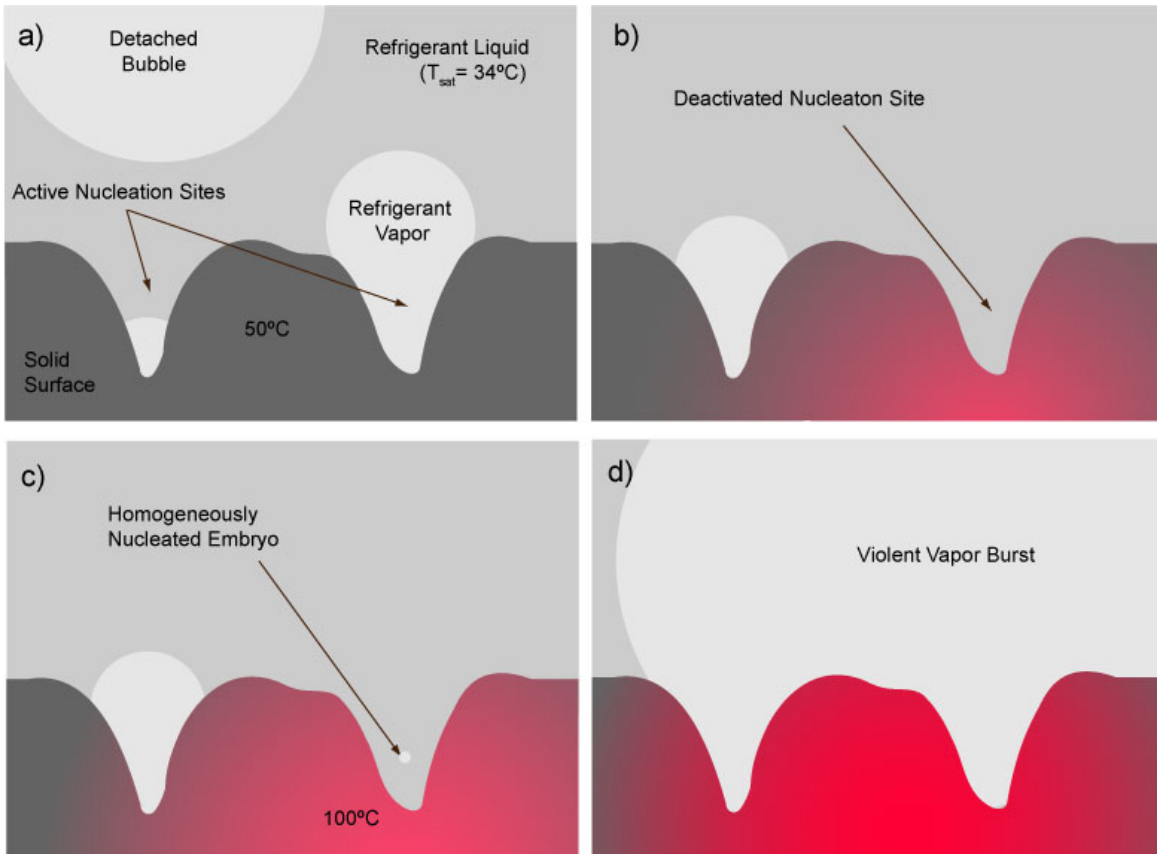


Fig. 4.3 – Sequence of site reactivation *without* AIBN surface treatment.

IV.3.2 Site Reactivation with AIBN-based Self-Nucleation

Just as before, deactivation of the nucleation site reduces the local heat transfer coefficient, which leads to an increase in wall temperature and the formation of an underlying hot spot (Fig. 4.4b). However, as this hot spot grows, it also heats up the AIBN surface directly beneath the deactivated site. Once the AIBN reaches its thermal decomposition temperature (approximately 64°C), it begins to release gaseous nitrogen, which then fills the cavity and creates an “artificial” entrapped bubble (Fig. 4.4c). This “injected” bubble provides a gas-liquid interface that allows further vaporization of the refrigerant to occur without the excessive superheat required to (homogeneously) nucleate a new embryo. As a result, *we prevent the large temperature gradients and explosive vapor burst associated with the untreated site*. Furthermore, controlled vaporization at the gas/vapor-liquid interface allows for a return to the normal ebullition cycle, which, in turn, raises the local heat transfer coefficient back to the original stable nucleate boiling level. With the original heat transfer coefficient restored, the localized hot spot disappears, the wall temperature decreases below 64°C , and the nitrogen effluence ceases (Fig. 4.4d). Thus the entire process is passive and self-regulating.

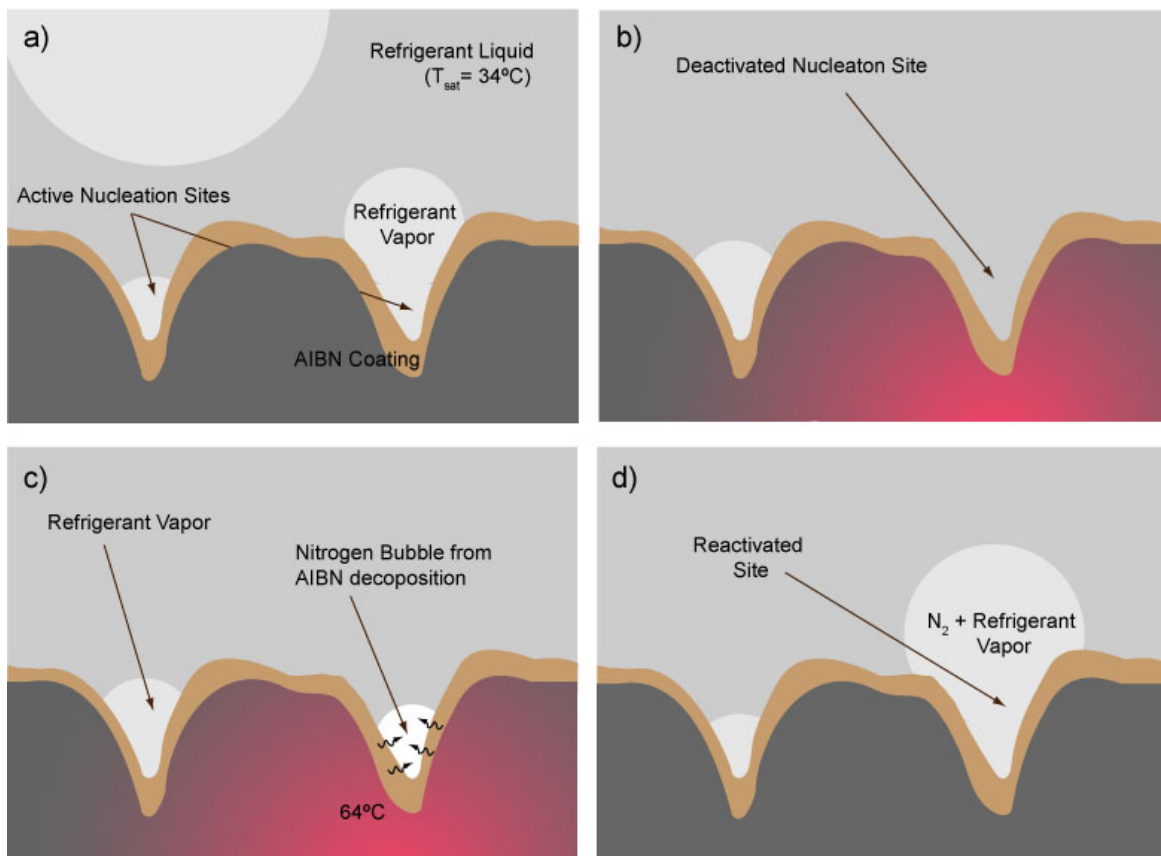


Fig. 4.4 – Sequence of site reactivation *with* AIBN surface treatment.

IV.4 Concept Limitations and Extensions

At this point I wish to examine some of the limitations of this concept, as well as some extensions and modifications to deal with these limitations. First, it must be stated that the thermal decomposition reaction is “one-way”, and therefore the nitrogen within the AIBN-polymer matrix will eventually be consumed. However, one must keep in mind that site reactivation via nitrogen effluence is expected to be a “rare” event. During normal steady-state operation, most sites will *not* become deactivated, and therefore the self-nucleation concept is intended to be more of a “safety valve” for a usually stable system than a conventional nucleation method. Furthermore, even a microscale evaporator will likely have hundreds of potential nucleation sites. If the nitrogen is completely exhausted beneath one site, there are countless nearby sites to dissipate the thermal load. In addition, because the thermal decomposition occurs gradually from 65-100°C, the excess surface temperature required for site reactivation will also only increase *gradually*. For example, assume again that the ideal wall temperature during operation is 50°C. A deactivated site will produce a localized hot spot of approximately 65-70°C before the underlying AIBN decomposes and injects nitrogen to reactivate the site. If that *same* site becomes deactivated again, the subsequent hot spot will probably reach 70-75°C before the decomposition can continue. The next time would be 75-80°C, and so on. Not only is this far superior to the homogeneous nucleation of untreated sites, it also provides some means of monitoring the remaining nitrogen levels in the AIBN-polymer matrix. For example, if surface temperatures reach 100°C before nucleation occurs, the thermal decomposition reaction has essentially concluded and the nitrogen supply has been exhausted.

A second drawback of my system in particular (but not the self-nucleation concept in general) is the temperature requirements imposed by the thermal decomposition of AIBN. Clearly a system with a nominal operating wall temperature of 100°C will not be compatible with AIBN. The same can be said of a system whose maximum wall temperature must not exceed 60°C. However, self-nucleation at higher or lower temperatures is still feasible using different chemicals. Each azonitrile, for example, has its own unique decomposition temperature, and it certainly goes without saying that many other chemicals produce insoluble gases upon thermal decomposition.

Finally, it should be noted that self-nucleation is not limited to thermally activated decomposition reactions. For certain refrigerants (*e.g.* water), electrolysis might be a viable option. A more complicated system might involve thermally-*triggered* injection, rather than thermally-*actuated* injection. For example, the ceramic known as yttria-stabilized zirconia (YSZ) conducts oxygen ions under external electrical bias. If the bias can be turned on or off based on the surface temperature (measured using thin-film resistors, for example), oxygen could be injected into cavities when the surface temperature indicates a departure from stable nucleate boiling. In addition, this oxygen could also be gettered in the “condenser” region of the device and returned to the YSZ to create a continual, non-exhausting system.

Chapter V. Experimental Design

V.1 Motivation and General Approach

The goal of this experimental study is to determine whether an evaporator surface treated with an AIBN coating can significantly lower the characteristic superheat required to initiate nucleation during cyclic pool boiling. Although based on microscale phenomena, the experimental study itself was purely macroscale in nature, primarily to limit both cost and fabrication time (*i.e.* to eliminate dependence on the Microlab, for which I am not yet qualified). Moreover, the principal purpose of this study is to serve as a mere proof-of-concept for self-nucleation, providing an additional avenue for future enhancement of the TGP program.

The general experimental approach is quite simple. Take two surfaces with essentially identical physical characteristics (thermal conductivity, wetting properties, surface roughness, *etc.*), but beneath one surface, embed a thin layer of encapsulated AIBN. Next, expose both surfaces to a large pool of saturated refrigerant with a compatible boiling point, and then heat both surfaces from below with gradually increasing flux. As the heat flux is cyclically increased past the onset of nucleation and then decreased until nucleation ceases, continuously measure the temperature of both surfaces and look for variations in the maximum surface temperature between the AIBN and no-AIBN sides. In theory, deactivated sites with underlying AIBN will require much lower superheat to reinitiate nucleate boiling because nitrogen gas released into cavities will provide a gas-liquid interface for low-superheat vaporization without requiring homogeneous embryo formation. Such variations may not be seen during *every* heating cycle, of course, since only when a sufficient number of sites have become deactivated will there be a significant temperature rise over a measurable area of the surface. This is especially true during the initial cycles, when insoluble gases remain entrapped in cavities. At some point, however, the large superheat disparity between heterogeneous self-nucleation and traditional homogeneous nucleation should be reflected in the cyclic temperature maxima (see Fig. 5.1).




Fig. 5.1 – Expected qualitative results of cyclic boiling experimental study.

V.2 Experimental Apparatus

V.2.1 Double Boiler and Auxiliary System Control

A schematic of the entire experimental apparatus is shown in Fig. 5.2, while an actual photograph of the corresponding system is shown in Fig. 5.3. In essence, the entire setup functions as a “double boiler” with several auxiliary control systems. The outer boiler is a 16 L (45.7 x 25.4 x 19 cm) open acrylic bath distributed by the PolyScience company. This outer boiler is filled with DI water and regulates the bulk temperature of the refrigerant in the cylindrical inner boiler. Temperature control for the outer boiler is provided by a 115V Haake C10 immersion circulator manufactured by Thermo Scientific. It has a 1000-watt heater, temperature accuracy of $\pm 0.04^{\circ}\text{C}$, and a circulation rate of 17 L/min.

The inner boiler consists of a 14” long section of 5" OD, 1/8" thick wall polycarbonate pipe bonded to a round quarter-inch thick polycarbonate bottom plate. Due to the extreme physical properties of the refrigerants (see Section V.4 below), the standard polycarbonate bond (SC-325 Polycarbonate Cement from U.S. Plastics) was reinforced along the interior seam with E-6000 industrial sealant. The cylinder is capped with an enamel-sealed wood lid, with two through-holes for the condenser plumbing. Both the inner boiler and immersion heater are mounted to a 3/4" plywood lid that fits over the outer boiler. The inner boiler can be quickly removed for filling and draining by loosening the three hex-head bolts, which screw through aluminum L-brackets into shallow (non-through) holes in the exterior wall of the inner boiler.

To prevent excessive refrigerant loss by evaporation, the inner boiler contains a 3" diameter condenser coil made from 1/4" OD copper tubing, which circulates 0°C water from an exterior ice water reservoir. Circulation is accomplished by a Hydor Seltz L20 II aquarium pump submerged in the ice water reservoir. The 115V 14W pump has a maximum circulation of 185 gal/hr, and the intake is protected by a wire-mesh screen to prevent small ice cubes from blocking the intake or entering the pump. Because pressures throughout the chilled water path remain low (and minor leakage is tolerable), the vinyl hoses are connected to the flanged copper tube condenser ports using simple stainless steel wire ties. The 1/2"-to-1/4" tubing transition is made using standard brass nipple fittings with threads sealed with Teflon tape.




Fig. 5.2 – Schematic of double boiler experimental apparatus.




Fig. 5.3 – Photograph of entire experimental setup.

V.2.2 Submersible Test Platform

The most critical elements of the experimental setup are contained within a submersible test platform that essentially functions as a small “diving bell”. A schematic pseudo-cross-section and photograph are shown in Fig. 5.4. The bell, or housing, is a 2-5/16" (5.874 cm) square injection-molded butyrate shell, whose top outer surface is machined and patterned with AIBN-polymer (see Section V.3 for a detailed description of this process). During testing, the entire housing is lowered into the refrigerant-filled inner boiler, and the top outer surface of the housing serves as the actual boiling surface. The one-piece injection-molded diving bell design prevents the highly-wetting refrigerant from penetrating the housing and contacting the various internal components. A 767 g steel base plate serves as ballast to keep the otherwise buoyant air-filled housing fully submerged, and regular flat washers serve as spacers between the base plate and housing assembly to provide clearance for the control and data wires.




Fig. 5.4 – Photograph and schematic pseudo-cross-section of submersible test platform.

Inside the butyrate housing is a custom-machined two-piece aluminum heater block that, in turn, encases a small (19 mm x 19 mm x 2.5 mm) ceramic resistive heater (see Fig. 5.5 for photographs and dimensioned schematic). The Ultramic 600 heater, manufactured by Watlow Electric and distributed by Therm-X of California, is fabricated from aluminum nitride and has a maximum power output of 16 W/cm² at 120V, a maximum operating temperature of 600°C, and contains an integrated K-type thermocouple. To minimize contact resistance, the ceramic heater is first coated with a thin layer of Radio Shack silicone-base heat sink compound before final encasement in the heater block.





Fig. 5.5 – Photographs and dimensioned schematic of the inner surface of the heater block.


The aluminum heater block assembly has several other key features. A raised boss on the top side is divided into two halves; one side will supply heat to the butyrate surface treated with AIBN, the other side to the untreated surface. The $1/8"$ air gap between the bosses inhibits lateral conduction and thus facilitates comparative measurements of AIBN performance. On the polished top surface of each boss is a series of finely machined pockets and grooves (see Fig. 5.6 for photo and dimensioned schematic), which contain precision thermistors embedded in conductive polysynthetic silver paste (Arctic Silver 5) and covered with a protective layer of aluminum tape. The commercial-grade leaded NTC thermistors (model QTMB-14C3, manufactured by Quality Thermistor) have a maximum bead diameter of $0.038"$ for extremely fast thermal response ($\tau_{\max} = 1$ sec, dissipation constant = 1 mW/ $^{\circ}\text{C}$, both in still air). They have a nominal resistance of 10 k Ω at 25°C , an operating range of -55 - 125°C , and a tolerance of $\pm 0.2^{\circ}\text{C}$ from 0 - 70°C .



Symmetric along $x = 0$

Any apparent square symmetry
is indeed symmetric

Milled pockets use 5/64" cutter
at 0.03" depth



Thermistor lead grooves use 60°
carbide engraver at 0.015" depth

All dimensions in inches

For location of thru holes, please
see schematic of interior surface

Fig. 5.6 – Photograph and dimensioned schematic of the top surface of the heater block.

The heater block is spring-loaded against the bottom side of the butyrate housing's top surface to normalize the compression factor in the contact resistance. Silicone-base heat sink compound is also applied at this interface to minimize contact resistance. The thermistor leads are fed through small holes along the perimeter of the heater block and held firmly in place by #2 set screws to prevent them from pulling out or failing by fatigue. The leads are then soldered to a 9-pin D-Sub connector, which is mounted on a custom-machined aluminum H-shaped bracket (which can also support an additional D-Sub connector if more sensors are desired). Below the H-shaped bracket is another milled-out aluminum block with four small laterally-drilled holes, into which set pins are inserted (through the sidewalls of the butyrate housing) to hold the housing in place. The D-Sub connector is then interfaced with ribbon cable via crimp-type connectors to transmit the sensor signals to the data acquisition system (see Section V.2.3 for further details on data acquisition). In addition, the power leads of the ceramic heater are wired to a Hewlett Packard E3630A DC power supply, which allows control over the heat flux through the boiler surface, and the K-type thermocouple leads are also wired to the data acquisition system.

V.2.3 Data Acquisition and Heater Control

Data acquisition and heater control is accomplished using a high-performance USB-based multifunction DAQ from National Instruments (model 6221, see Fig. 5.7) utilizing a custom virtual instrument constructed using Labview.




Fig. 5.7 – Data Acquisition System (National Instruments model 6221) with circuits and connections labeled.

A flowchart of the data stream and control is shown in Fig. 5.8 and can be explained simply as follows: A 5V bias is applied in sequence across each of the embedded thermistors. A series of Wheatstone bridge circuits (see Fig. 5.10) are used to perform comparative readings to accurately measure the voltage drop across each thermistor (subscript th). This voltage drop is then converted into a corresponding resistance by the following relations:

$$V_{th} = \frac{R_1}{R_1 + R_4} V_0 - \frac{R_{th}}{R_S + R_3} V_0 = \frac{\frac{R_1}{R_4} - \frac{R_{th}}{R_3}}{\left(1 + \frac{R_1}{R_4}\right) \left(1 + \frac{R_{th}}{R_3}\right)} V_0 \quad (29)$$

Letting $R_1 = R_3 = R_4 = R$:

$$V_{th} = \frac{1 - \frac{R_{th}}{R}}{2 \left(1 + \frac{R_{th}}{R} \right)} V_0 \quad (30)$$

Rearranging and solving for R_{th} :

$$R_{th} = \frac{V_0 - 2V_{th}}{2V_{th} + V_0} R \quad (31)$$

Sensitivity is maximized by using bridge resistors that correspond to the nominal resistance of the thermistors ($R = 10 \text{ k}\Omega$). Using the standard Z-curve coefficients, this measured resistance is converted to its corresponding temperature using the relation below, which is then recorded on a strip chart.

$$\frac{1}{T} = A + B \ln R + C(\ln R)^2 + D(\ln R)^3 \quad (32)$$

$$A = 0.001116401465500$$

$$B = 0.000237982973213$$

$$C = -0.000000372283234$$

$$D = 0.000000099063233$$

Additional controls on the front panel of the virtual instrument control the ceramic heater by specifying the heating time and total cycle time (see Fig. 5.9 for a screen shot of front panel). A digital output from the DAQ is connected to the gate of a 60W IRF530 amplifier IC, while the source and drain terminals connect the power supply to the ceramic heater (see Fig. 5.11).

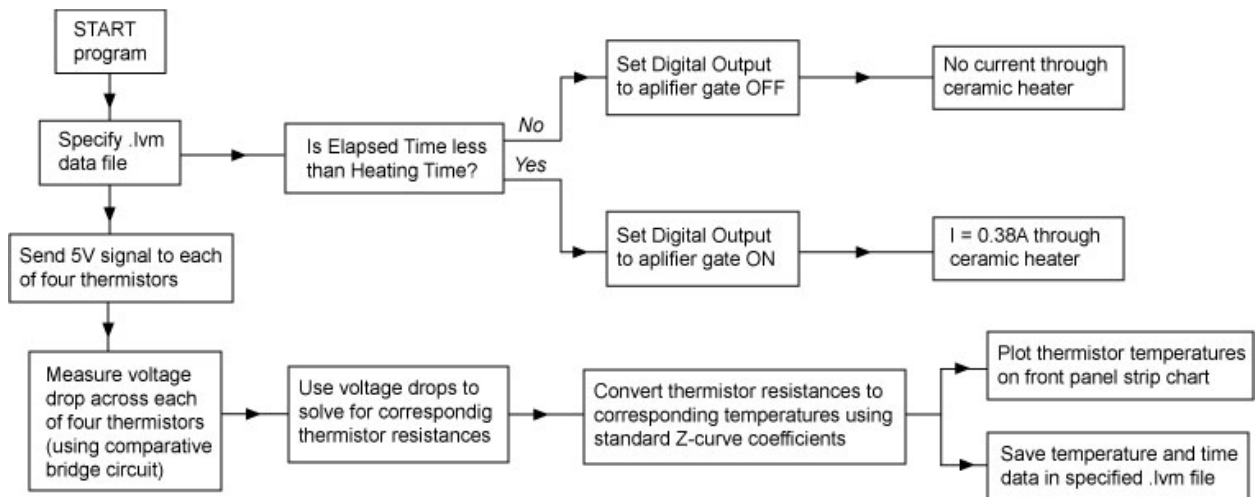


Fig. 5.8 – Simplified flowchart of basic Labview program control.

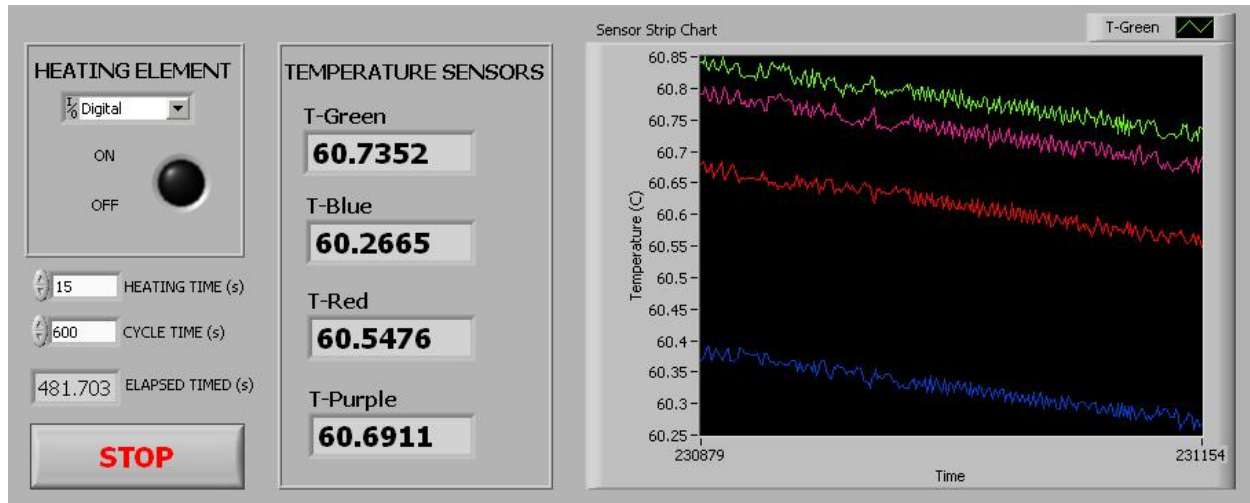


Fig. 5.9 – Screenshot of Labview Virtual Instrument front panel.

A red LED is also wired to the amplifier for visual confirmation of current flow when the heater is on. For this study, heater control amounts to simple on/off control (power supply is permanently set to 23.8V at 0.38A). For the majority of my experiments, this means that the heating cycle is set such that the heater is powered on long enough to initiate nucleate boiling on the surface and then powered off long enough to fully deactivate nucleate boiling.




Fig. 5.10 – Standard Wheatstone bridge circuits for precise measurements of thermistor voltage drops.




Fig. 5.11 – Amplifier circuit for controlling and powering ceramic heater.

V.3 Test Surfaces

The test surfaces are CNC-milled on the top of the outer surface of the butyrate housings. The machined pattern consists of an array of circular shallow wells superimposed on two deep oblong wells, which are centered directly ovetop (but completely *within*) the “footprint” of each of the two halves of the split heater boss (see Fig. 5.12 for photograph and milling parameters). The butyrate housing walls are 0.055" thick, leaving a membrane thickness of 0.02" under the shallow wells after machining.

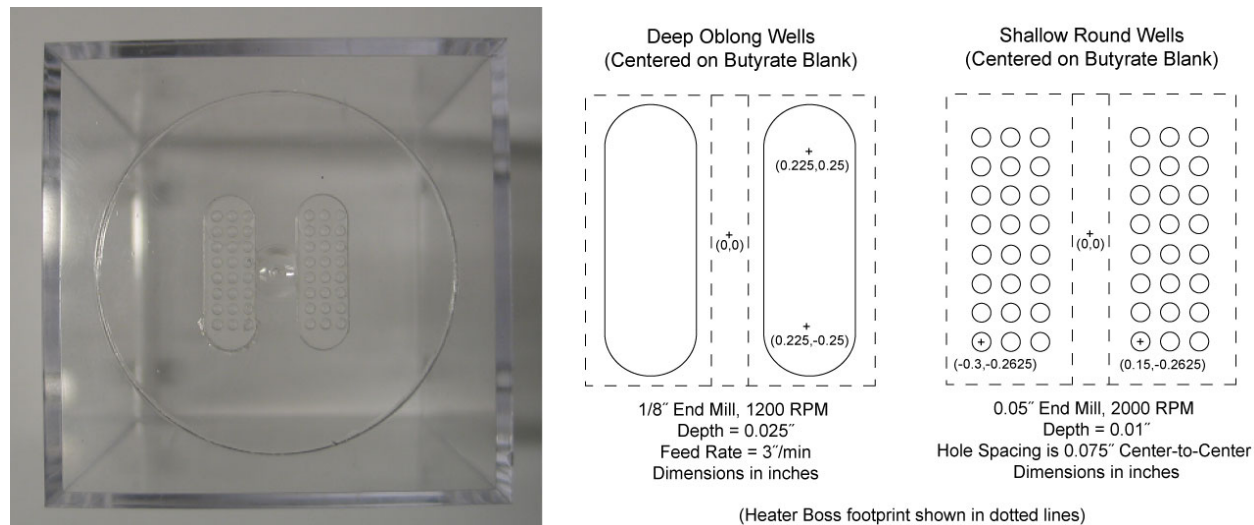


Fig. 5.12 – Photograph of machined butyrate housing prior to AIBN-polymer coating, as well as CNC milling dimensions and parameters.

This “double well” pattern was a compromise to satisfy two important design constraints. First, it is absolutely critical to retain overall structural rigidity of the wall so that the spring-loaded heater block maintains uniform contact pressure with the butyrate wall. At the same time, however, it is desirable to have as thin a butyrate membrane as possible so that thermal resistance across the butyrate is minimized and thermistor temperature data is most characteristic of the actual surface temperature. For this reason, the deep oblong wells must lie entirely within the footprint of the heater boss or else significant wall deflection will create uneven surface contact (see Fig. 5.13).

The 3x8 array of 0.01"-deep shallow wells hold the AIBN-polymer mixture, which is applied much like spackling in a screen-printing-like process. An array of small wells (as opposed to a single large well) helps minimize the shear and liftoff of AIBN-polymer mixture during screen-printing and also provides a more “statistically uniform” surface from one side of the test surface to the other. Finally, four through-holes are bored laterally through the sides of the housings to allow inserted set pins to constrain the housing over the spring-loaded heater block assembly.

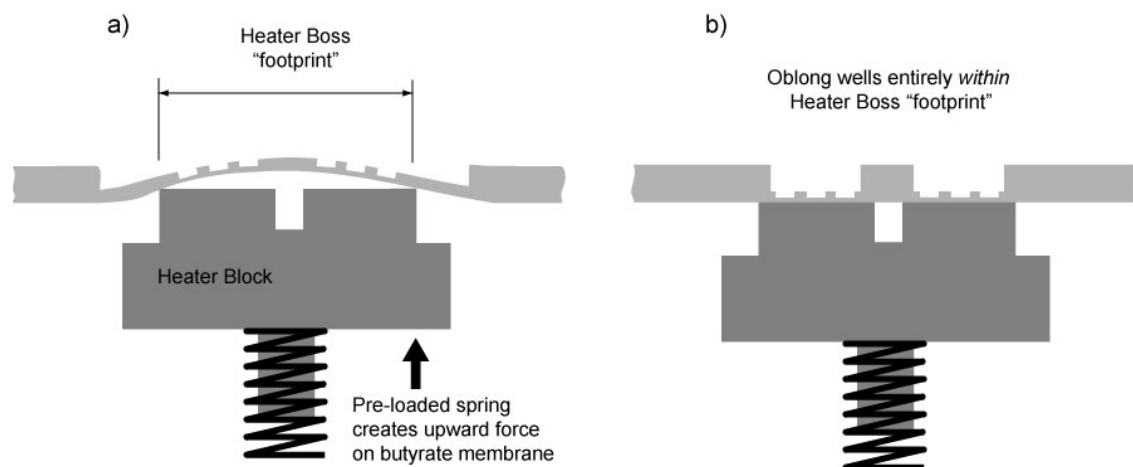


Fig. 5.13 – a) If the butyrate surface is milled too thinly *outside the heater boss footprint*, the heater block will deflect the membrane and create uneven surface contact. **b)** The double well design (with thin milling only completely *within the heater boss footprint*, even surface contact is maintained.

After machining, the butyrate housings are first cleaned with mild soap and water before the surface treatments are applied (see Fig. 5.14 for process flow). On one side, the shallow wells are filled with a mixture of AIBN and polymer (mix ratio is 150 mg AIBN to 0.2 mL of Zinsser BIN shellac-base liquid primer). The shallow wells on the other side are filled with primer alone. After the excess primer/mix has been carefully scraped off, a thin layer of Zinsser BIN shellac-base aerosol primer is applied to both sides to ensure that both sides have identical physical surface characteristics. A completed sample is shown in Fig. 5.15.




Fig. 5.14 – Process flow for butyrate housings: a) Milled surface is cleaned; **b)** AIBN and liquid primer mix is applied to treated side, liquid primer alone to untreated side; **c)** Excess primer and/or AIBN is scraped off of both sides; **d)** Light coat of aerosol primer is applied to both sides.



Fig. 5.15 – Completed housing sample ready for boiling tests.

V.4 Refrigerants

Two different refrigerants were used in this study, HFE 7000 and HFE 7100. Both are Novec Engineered Fluids produced by the 3M company consisting of fully fluorinated hydrocarbon chains, with HFE 7100 possessing one additional fluorinated carbon atom. The first set of experiments involved HFE 7000 ($T_{sat} = 34^{\circ}\text{C}$), but unexpected boiling characteristics required a shift to a refrigerant with a significantly higher boiling point (see Chapter VI). Consequently, the second set of experiments involved HFE 7100 ($T_{sat} = 61^{\circ}\text{C}$).

V.4.1 Novec HFE 7000

Novec HFE 7000 (scientific name is methoxy-heptafluoropropane) is a clear, odorless, volatile, nonflammable fluid with the following relevant physical properties (at 25°C and 1 atm unless otherwise noted):

Chemical formula:	$\text{C}_3\text{F}_7\text{OCH}_3$
Molecular weight:	200 g/mol
Boiling point:	34°C
Liquid density:	1400 kg/m^3
Kinematic viscosity:	0.32 cSt
Surface tension:	12.4 mN/m
Latent heat of vaporization:	142 kJ/kg
Specific heat:	1300 J/kg·K
Vapor pressure:	64.6 kPa
Solubility of air in fluid:	~35 vol %
Dielectric strength:	~40 kV

(Source: 3M)

There are several properties that make this fluid extremely well-suited to this experimental study. Most notably is its boiling point. As discussed in Section IV.3, a boiling point of 34°C is compatible with AIBN's thermal decomposition temperature of 64°C. Based on literature from 3M, boiling incipience should occur at a wall temperature around 60-65°C, while stable nucleate boiling should be maintained at a wall temperature around 50°C; thus self-nucleation via thermal decomposition should occur only when localized hot spots develop under deactivated sites. In addition, the dielectric properties ensure that any leakage will not short electrical signals, and the combination of low surface tension (compare to 72.8 mN/m for water) and air solubility will test the extreme limits of cavity wetting and site deactivation.

In addition to these favorable characteristics, there are several aspects that make HFE 7000 somewhat difficult to work with. Although low surface tension is beneficial from an experimental standpoint, it becomes a serious design issue in terms of seals, particularly when combined with extraordinarily low viscosity (compare to 1.0 cSt for water). Seals that will normally hold water at 10 atm will leak HFE 7000 at 1 atm. Consequently, the polycarbonate-polycarbonate seals for the inner boiler had to be reinforced with industrial sealant, and the "diving bell" housing had to be made from a single injection-molded piece. Earlier attempts at sealed two-piece housings failed miserably. The near-room-temperature boiling point also implies high vapor pressure (compare to 3.2 kPa for water). This necessitates the copper coil condenser to prevent excessive fluid loss. One obvious shortcoming of HFE 7000 as a phase-change refrigerant is its low latent heat of vaporization (again, compared to water, which at 2260 kJ/kg, has the second highest latent heat of any known fluid).

V.4.2 Novec HFE 7100

Novec HFE 7100 (scientific name is methoxy-nonafluorobutane) is also a clear, odorless, volatile, nonflammable fluid with the following relevant physical properties (at 25°C and 1 atm unless otherwise noted):

Chemical formula:	C ₄ F ₉ OCH ₃
Molecular weight:	250 g/mol
Boiling point:	61°C
Liquid density:	1520 kg/m ³
Kinematic viscosity:	0.37 cSt
Surface tension:	13.6 mN/m
Latent heat of vaporization:	111.6 kJ/kg
Specific heat:	1180 J/kg·K
Vapor pressure:	26.7 kPa
Solubility of air in fluid:	53 vol %
Dielectric strength:	28 kV

(Source: 3M)

Based on their molecular similarity and identical functional groups, it is not surprising that HFE 7100 is quite comparable to HFE 7000 in terms of wettability and vapor pressure. One significant difference, though, is a much higher boiling point. While in theory, the boiling point

of HFE 7000 makes it ideally suited to AIBN-based self-nucleation experiments, unexpected results necessitated a shift to HFE 7100 (see Chapter VI).

V.5 Experimental Controls and Thermal Considerations

The crudeness, simplicity, and low cost of this macroscale proof-of-concept study cannot be overemphasized. The difficulty and complexity associated with fabricating and collecting data from individual nucleation sites with highly-controlled parameters has been largely circumvented by exploiting a few clever design tricks and by invoking a few reasonable thermal assumptions.

First, this experiment seeks primarily to simply compare the maximum measured surface temperature achieved during the boiling cycle on surfaces that either are or are not treated with AIBN. This temperature is a direct indicator of the superheat required to initiate nucleate boiling and thus provides a means of testing the self-nucleation concept. Consequently, extremely accurate temperature measurement and estimation of contact resistances are not necessary for a number of reasons:

1. The concept (in theory) should produce large disparities in temperature (on the order of tens of degrees) between self-nucleated and untreated surfaces, even allowing for a great deal of attenuation due to lateral conduction (thermal leakage) and the short time scales associated with nucleation events.
2. The experiment will examine a large number of boiling cycles to eliminate any of the stochastic variation associated with any one particular nucleation event. While it is not expected that *every* cycle will result in many high-superheat homogeneous nucleation events on the untreated side, some cycles (when many nucleation sites have become deactivated) should produce an easily measurable temperature disparity.
3. Higher surface temperatures will be reflected in higher thermistor temperatures, regardless of small variations in membrane thickness, contact resistance, or aerosol coating from one trial to the next. Again, this setup is designed to emphasize *differences* between the “AIBN” and “no-AIBN” sides of the heated boss, rather than explicitly determining boiling parameters like superheat and heat flux.

Second, since significant lateral conduction is expected *on each side of the divided boss* (but not between), self-nucleation at just *some* of the AIBN-filled wells should be reflected in an overall lower maximum temperature recorded by the “AIBN-side” thermistors. This precludes the requirement for microscale temperature sensors beneath every single nucleation site. The poor thermal conductivity of the plastic membrane prevents lateral conduction from the no-AIBN to the AIBN side, and the air gap in the divided boss was designed with the goal of minimizing lateral conduction across the aluminum heater block. Lastly, two thermistors were embedded on each side of the heater boss to confirm that there was not radically uneven contact pressure or uneven heating (or in case a thermistor failed). For conclusions regarding the relative success of these design philosophies and assumptions, please see Chapter VII.

Chapter VI. Experimental Results and Analysis

VI.1 Encapsulation and Thermal Decomposition Test

A preliminary test was required in this study to confirm the functionality of the encapsulation technique. It had to be shown that AIBN could be successfully incorporated into a stable polymer matrix that would not dissolve in the refrigerant or flake off during heating. In addition, it had to be shown that effluent nitrogen gas could diffuse through the matrix and escape the surface once the AIBN was heated to its thermal decomposition temperature.

First, the shallow well pattern (see Section V.3) was CNC-milled into a small sample of 1/4" acrylic. The deep oblong wells were not required for this test because the sample was heated externally (by immersion in hot water), rather than by conduction through the plastic itself. The AIBN and polymer mixes were then applied using the same process depicted in Fig. 5.14. After curing, the sample was submerged in a Pyrex beaker of DI water and placed on a hot plate. The water was then heated above the thermal decomposition temperature of AIBN (but well below the boiling point of water). Bubbles were clearly visible emerging from the surface of the wells treated with AIBN, while no bubbles appeared on the untreated side (see Fig. 6.1)




Fig. 6.1 – Nitrogen effluence from treated surface during preliminary tests in heated water.

VI.2 Cyclic Boiling Tests with HFE 7000 Refrigerant

Cyclic boiling tests with HFE 7000 ($T_{sat} = 34^{\circ}\text{C}$) were intended to mimic a reasonably realistic application scenario. Namely, it was expected that initial heterogeneous nucleation from entrapped gases during the first boiling cycle (or during prolonged, stable ebullition) would occur between 50-60°C and thus not trigger nitrogen effluence, but that subsequent cycles (after site deactivation) would increase superheat requirements into the range of AIBN thermal

decomposition. Such self-nucleation would prevent the AIBN-treated side from exceeding approximately 65-70°C, while the untreated side might reach much higher temperatures before nucleation occurred.

VI.2.1 Saturated HFE 7000

The first set of boiling experiments was performed with saturated HFE 7000 refrigerant. The water bath was heated to 36°C to maintain an inner boiler temperature of 34°C. For this test, the heater control was set to turn on for 100 sec and then turn off for 200 sec. As discussed in Section V.2.3, this ratio allowed sufficient heating to initiate nucleate boiling on both test surfaces and sufficient cooling to cease nucleate boiling on both surfaces. In these first tests, I did not let the surfaces cool all the way back down to the bulk liquid temperature, so some trial and error was necessary to find a stable heating cycle. Once that was established, a total of 9 cycles were conducted. The rapidly rising bubble jets characteristic of all the pool boiling experiments in this study can be seen in Fig. 6.2. Plots of the temperature data for the four thermistors (two below each side) are shown in Fig. 6.3. Only one curve is visible because the temperature curves for all four thermistors are so close that the curves lie on top one another in this zoomed-out view. Fig. 6.4 shows a zoomed-in view of one of the peaks, so all four signals can be seen.




Fig. 6.2 – Characteristic faint, rapidly rising bubble jets associated with all pool boiling experiments in this study.




Fig. 6.3 – First cyclic boiling test with saturated HFE 7000 ($T_{bulk} = 34^{\circ}\text{C}$).




Fig. 6.4 – Zoomed-in view of one of the peaks in Fig. 6.2. The red and blue curves are the AIBN-treated sides. The pink and green curves are the untreated side.

What is immediately striking about this data is the extremely low superheat required for the onset of nucleate boiling. Nucleation occurred on both surfaces at approximately 48°C, well below the thermal decomposition temperature of AIBN, for all cycles. Moreover, keep in mind that the true surface temperature is inherently cooler than that measured by the thermistors, which lie below an additional 0.01" of AIBN-polymer and 0.02" of butyrate (both of which have poor thermal conductivity) and across a finite contact resistance between the heater block and the butyrate housing. This was truly a surprising result. When selecting a suitable refrigerant, I had relied on characteristic boiling curves for Novec fluids provided by 3M, which predict superheats of ~40°C (see Fig. 6.5).

This same test was repeated several times but always produced similar results. My initial conclusion was that the test surfaces in this study must have much higher surface roughness or cavity density than the surfaces used in 3M's tests. Certainly this is a reasonable possibility given the likelihood of their using highly polished metal test surfaces more likely to resemble those found in traditional industrial heat transfer applications for which the fluid was engineered. Supposing this conclusion was indeed correct, perhaps it was also possible that nucleation sites on either side had simply not yet become deactivated. As discussed in Section V.4.1, such low-viscosity, low-surface-tension fluids were expected to be *excessively* wetting, perhaps even flooding many nucleation sites during initial submersion of the test platform. If not, however, the next logical step was to leave the test platform immersed for an extended time prior to initial heating, as Xu and Lloyd had done in [4]. Before the next saturated cyclic boiling test, the test platform was immersed for 72 hours at approximately 25°C (~10°C of subcooling), but the effects on required superheat were negligible. Nucleate boiling was still observed on both test surfaces well below the thermal decomposition temperature of AIBN.




Fig. 6.5 – Characteristic boiling curve data for Novec fluids provided by 3M.

While the results were unexpected and somewhat disappointing, two significant pieces of information *were* gained from the first round of experiments. The first, while neither revolutionary nor unexpected, was that the AIBN treatment does not affect the heat transfer

characteristics off the surface *below* the thermal decomposition temperature. In other words, since all four sensors (on both sides) recorded essentially identical temperatures throughout the tests, we can confirm that the AIBN-treated surface does not have some radically different thermal conductivity or disparate surface characteristic. Thus any differences seen in later tests can be more confidently attributed to nitrogen effluence and not an intrinsic difference between the surfaces themselves.

The second useful observation was visual confirmation of the expected asymmetry between the temperature required for the onset of nucleate boiling during heating and the temperature at which nucleate boiling ceased during cooling (see Section II.3.3). In this set of experiments, for example, nucleation occurred on both sides at approximately 48°C during heating but persisted during cooling until temperatures dropped below approximately 41°C. This phenomenon was always observed to some degree in every test throughout the study.

VI.2.2 Subcooled HFE 7000

Since prolonged subcooled immersion *prior to saturated heating* did not increase the nucleation temperature into the range where the effects of AIBN decomposition could be observed, the next logical step was to perform similar cyclic boiling experiments with subcooled HFE 7000. A 60-hour immersion at 25°C (~10°C of subcooling) prior to heating was also performed before this set of tests, as well. After this time, the inner boiler bulk temperature was raised to 28°C (6°C of subcooling), and the heater control was set to turn on for 150 sec and then turn off for 100 sec. Notice that with subcooling, the required heating time to initiate nucleation is longer than for saturated boiling, while the cooling time for nucleation cessation is shorter. This is, of course, expected, since the driving force (temperature difference) between the wall and bulk liquid is larger, and thus heat is more easily transferred. A total of 5 cycles were conducted initially, and plots of the temperature data are shown in Fig. 6.6.

In this case, *thermistor* temperatures definitely reached the thermal decomposition temperature of AIBN prior to nucleation (which occurred at approximately 75 °C on both sides). As the graphs indicate, however, there was still negligible difference in the temperatures recorded by all four thermistors during all cycles, indicating no measurable effect due to self-nucleation. Again, this was a frustrating result, but there remained the very good possibility that thermal resistance between the thermistors and the actual surface had prevented the underlying AIBN from exceeding 64°C *prior to nucleation*. In other words, although the maximum temperature achieved during each cycle would certainly result in significant thermal decomposition and nitrogen effluence, perhaps this occurred *only after the surfaces had already nucleated by traditional means*. Moreover, keep in mind that 64°C is only the *initial* decomposition temperature. The rate of nitrogen effluence increases with temperature, particularly above 80°C (see Fig. 4.2b). Therefore, it was still reasonable to suppose that a surface-fluid combination with still higher nominal nucleation temperature might produce positive results. Since control over the surface was extremely limited by the fabrication techniques, it was decided to switch to a refrigerant with a higher boiling point.




Fig. 6.6 – Cyclic boiling test with subcooled HFE 7000 ($T_{bulk} = 28^{\circ}\text{C}$).

VI.3 Cyclic Boiling Tests with HFE 7100 Refrigerant

Because the boiling point of HFE 7100 ($T_{sat} = 61^{\circ}\text{C}$) is so close to the thermal decomposition temperature of AIBN, these tests were *not* intended to mimic a realistic application scenario. This is because initial heterogeneous nucleation due to entrapped gases (or even prolonged, stable ebullition) would be expected to occur between $70\text{--}80^{\circ}\text{C}$, thus already exhausting much of the nitrogen supply during “normal” operation. However, at this point, I simply wanted to see if even higher homogeneous nucleation temperatures could be precluded due to nitrogen effluence between $80\text{--}100^{\circ}\text{C}$.

VI.3.1 Saturated HFE 7100

The first set of boiling experiments using HFE 7100 refrigerant were performed at saturation. The water bath was heated to 62.5°C to maintain an inner boiler temperature of 61°C . As discussed above, because the saturation temperature is only a few degrees below the thermal decomposition temperature of AIBN, I chose not to immerse the test platform for an extended period prior to surface heating because I did not want my initial nucleation event to completely exhaust the nitrogen supply. For this test, the heater control was set to turn on for 15 sec and then turn off for 585 sec. This long cooling time was necessary because nucleation persisted almost all the way down to the saturation temperature. A total of 20 cycles were conducted, and plots of the temperature data for the four thermistors are shown in Fig. 6.7.




Fig. 6.7 – Cyclic boiling test with saturated HFE 7100 ($T_{bulk} = 61^{\circ}\text{C}$).

Although slightly less so this time, it was still very surprising to observe nucleation at such low superheat (thermistors indicated approximately 64°C). Even using a refrigerant with a saturation temperature almost 30°C higher than HFE 7000, nucleation was still observed on both surfaces at temperatures below the thermal decomposition temperature of AIBN for all 20 cycles. Although the blue curve deviates slightly from the others (see Fig. 6.8 for a zoomed-in view), the deviation is less than 1°C and is much more likely due to problems with the thermistor itself than to the AIBN treatment. This is further confirmed by the fact that the red curve, also on the AIBN-treated side, was not affected and because the temperature difference in the blue curve is essentially constant and occurs at both the high and low temperature extremes.




Fig. 6.8 – Zoomed-in view of one cycle from Fig. 6.6.

VI.3.2 Subcooled HFE 7100

Since saturated boiling of HFE 7100 did not increase the nucleation temperature into the range where the effects of AIBN decomposition could be observed (and since prolonged immersion prior to heating did not have any effect during the tests with HFE 7000), the next logical step was to perform similar cyclic boiling experiments with subcooled HFE 7100. The inner boiler bulk temperature was lowered to 51°C (10°C of subcooling), and the heater control was set to turn on for 60 sec and then turn off for 180 sec. Again, note the longer heating times and shorter cooling times required for subcooled boiling. A total of 19 cycles were conducted, and plots of the temperature data are shown in Fig. 6.9, with a zoomed-in view of one cycle in Fig. 6.10. Once again, the blue curve deviates from the others (this time by approximately 2°C), but the difference is constant throughout the cycle and is most likely due to problems with the thermistor, which was subsequently replaced. With such extreme subcooling, nucleation did not occur until thermistor temperatures were approximately 80°C, so I know that the AIBN reached its thermal decomposition temperature prior to nucleation. However, during the initial trial and error calibration of the heating cycle, thermistor temperatures reached 95°C, so perhaps nearly all of the nitrogen had been exhausted prior to the cyclic test for which data was collected.




Fig. 6.9 – Cyclic boiling test with subcooled HFE 7100 ($T_{bulk} = 51^{\circ}\text{C}$).




Fig. 6.10 – Zoomed-in view of one cycle from Fig. 6.8.

VI.4 Cyclic Boiling Endurance Tests

One critical issue I had not yet explored fully was the possibility that the cyclic boiling tests I had performed thus far were simply not long enough in duration to show the effects of self-nucleation. In other words, perhaps they did not contain enough cycles to fully deactivate a large number of nucleation sites, either by carrying off insoluble gases during normal ebullition or otherwise. Consequently, I decided to run an additional test with each refrigerant for many more cycles.

VI.4.1 Subcooled HFE 7000 Endurance Test

Because the 6°C subcooled experiments with HFE 7000 cycled through maximum temperatures slightly higher than the ideal application scenario (where nominal heterogeneous nucleation would occur slightly below the thermal decomposition temperature), I decided to repeat boiling conditions similar to those of Section VI.2.2 but with slightly less subcooling. Bulk fluid temperature was 30°C (4°C of subcooling), heating time was 40 sec, and cooling time was 260 sec. With this degree of subcooling, normal nucleation occurred at approximately 50°C. The 13 hour, 20 minute test consisted of 160 total heating cycles, and the temperature data is shown in Fig. 6.11. Note that even with the large number of cycles, the maximum temperature of all thermistors remained well below the thermal decomposition temperature of AIBN. In other words, there were not enough high-superheat, homogeneous nucleation events to create a measurable hot spot above the thermistors. Consequently, it is not surprising that temperatures for all four thermistors tracked one another more or less exactly throughout the endurance test.

VI.4.2 Saturated HFE 7100 Endurance Test

At this point, I went back and reexamined the data for the saturated HFE 7100 test. Recall that I had initially assumed that the HFE 7100 tests would not be characteristic of an actual application scenario. However, because the nominal nucleation temperature (64°C) was so close to (yet still below) the initial thermal decomposition temperature (contrary to my prediction), this experiment, like the subcooled HFE 7000, actually presented a very realistic application scenario. Normal, heterogeneous nucleation (or steady ebullition) would, in fact, *not* drain the nitrogen supply; rather only nucleation events requiring higher-than-normal superheat would sufficiently raise the surface temperature to trigger AIBN decomposition. As a result of this finding, I decided it would be very worthwhile to run an endurance test with the heating cycle of the saturated test. Bulk fluid temperature was 61°C, heating time was 15 sec, and cooling time was 585 sec. The 9 hour, 10 minute test consisted of 55 total heating cycles, and the temperature data is shown in Fig. 6.12. As was the case with the HFE 7000 endurance test, the maximum temperature of all thermistors indicate that the surface temperature remained below the thermal decomposition temperature of AIBN. In other words, surface superheats never reached the point at which AIBN would be expected to have an effect on nucleation.




Fig. 6.11 – Subcooled HFE 7000 endurance test ($T_{bulk} = 30^{\circ}\text{C}$).




Fig. 6.12 – Saturated HFE 7100 endurance test ($T_{bulk} = 61^{\circ}\text{C}$).

VI.5 Quasi-static Temperature-Controlled Boiling

Up to this point, all experimental work had focused solely on attempting to observe differences in surface temperature (manifested to some qualitative degree in the underlying thermistor temperatures) while gradually increasing the heat flux to the nucleation point. And thus far, no significant temperature variation has been seen. Moreover, there was absolutely no slope change (let alone a temperature drop) at the onset of nucleate boiling. This caused me to step back and reconsider my basic assumptions. First, because the heater block has such a large thermal mass and such high thermal conductivity, it occurred to me that perhaps any attempt to control the heat flux *at the surface* was largely in vain. My entire temperature-based comparative method hinged on the assumption that heat flux through both surfaces would be equal. What if the heat flux had been asymmetrical? In other words, given a new heat transfer mode (*i.e.* nucleate boiling), the increased heat transfer coefficient *at the surface* could result in increased *flux*, rather than decreased *surface temperature* (or certainly decreased thermistor temperature). This is particularly likely with my particular experimental setup because the thermistors are embedded in the surface of the aluminum heater block, not at the boiling surface. Essentially, the entire surface of the heater block (and thus all four thermistors) could be isothermal even with much higher flux passing through one side! If this were indeed the case, then the situation is more like *temperature-controlled boiling* and not *flux-controlled boiling* (compare Figs. 2.2 and 2.4). This presents an entirely new method of observing self-nucleation: if all four thermistors indicate the same temperature, but only one side has commenced nucleate boiling while the other has not, then one surface has effectively undergone lower-superheat (or self-) nucleation.

To test this hypothesis, I needed first to ensure that the temperature of the heater block surface was indeed controlled during the onset of nucleation. This was not possible using my original heating method because the large thermal mass and high thermal conductivity of the heater block allowed surface temperatures to increase essentially linearly throughout the heating process, including during nucleation. (In fact, temperatures would actually continue to increase for 1-2 seconds *after* the power had been turned off.) To fix this, I adjusted the DC power supply to 0V, set the heating time equal to the cycle time on the virtual instrument front panel, and then *very slowly* increased the DC voltage, allowing sufficient time before each incremental increase in power for the block to reach steady state.

The second issue was selecting a suitable combination of refrigerant and subcooling level. Essentially, what I wanted to happen was for released nitrogen to initiate nucleation on the AIBN-treated side *before* normal nucleation occurred on the untreated side. Since AIBN decomposition does not begin until 65°C, I needed a combination that nucleated above 65°C. Looking back over previous cyclic boiling tests, I saw that HFE 7100 subcooled by 10°C did not nucleate until thermistor temperatures reached approximately 80°C. In addition, when reviewing the test parameters in my lab notebook, I noticed that during initial cycle calibration (for which I had not bothered to record data), I had observed that the AIBN-treated side had nucleated before the untreated side! Although I did not recognize the significance at the time, this now seemed to confirm my asymmetric flux hypothesis.

For the actual test, I again heated the bulk fluid in the inner boiler to 51°C (10°C of subcooling). Then I started the Labview program and began taking data as I incrementally raised the DC voltage from 0V. Between each increase, I waited until the maximum hash mark on the front panel real-time strip chart did not increase for 2 minutes. Overall, the ramp-up took approximately 4 hours. When the voltage had reached 11.24V, all four thermistors recorded steady temperatures ranging from 72.72°C to 73.12°C, indicating an essentially isothermal heater block. However, the AIBN-treated side had initiated and maintained steady nucleate boiling, while the untreated side did not! This is clearly visible in the series of photographs in Fig. 6.13. Nucleate boiling continued strictly on the AIBN-treated side until the voltage reached 12.49V, at which time the untreated side also initiated nucleate boiling. Steady-state thermistor temperatures ranged from 76.93°C to 77.35°C at this voltage. Such visual observations, in conjunction with the temperature at which they occurred, strongly support the possibility of self-nucleation. The complete temperature vs. time data is shown in Fig. 6.14, while Fig. 6.15 shows how the steady state temperature varied with supplied power. (Note: the solid lines between data points are merely linear connections.) Obviously, not all of this power is being transferred through the boiling surface; nevertheless, the total dissipated power should be at least *qualitatively* indicative of the surface heat flux, especially given my design efforts to thermally isolate the heater block as much as possible.




Fig. 6.13 – Visual evidence of self-nucleation during quasi-static temperature-controlled boiling.




Fig. 6.14 – Quasi-static temperature-controlled HFE 7100 test ($T_{bulk} = 51^{\circ}\text{C}$).




Fig. 6.15 – Quasi-static temperature-controlled HFE 7100 test ($T_{bulk} = 51^{\circ}\text{C}$).

Chapter VII. Conclusions

VII.1 Systemic Observations

The series of pool boil experiments conducted in this study displayed many characteristics supported by classical phase change heat transfer theory. For example, in all tests, the temperature required for the onset of nucleate boiling was higher than that required to sustain nucleate boiling once it had been initiated (see Table 7.1). Moreover, it was also observed that subcooling indeed raises the temperature required for the onset of nucleation and also raises the minimum temperature to which boiling can persist as the surface cools. While certainly valuable as experimental confirmation of academic theory, verification of these phenomena was not the focus of this study. What I had hoped to observe in the comparative temperature-based data was not seen at all. Namely, where I had expected to record dramatic differences in the temperature required for nucleation between the treated and untreated surfaces (see Fig. 5.1), I instead recorded essentially identical temperature curves throughout every boiling cycle. In addition, though the qualitative trends mentioned above gave me a great deal of confidence in the fidelity of my equipment and the aptness of my programming, they also highlighted major shortcomings in a few key elements that precluded the collection of meaningful quantitative data. Such shortcomings and their consequences are discussed in further detail in the next section.

Table. 7.1 – Summary of boiling parameters during cyclic boiling tests.

Refrigerant	Test Type	T _{bulk} (°C)	T _{bath} (°C)	Heating Time	Cooling Time	Number of Cycles	T _{ONB} (°C)	T _{DNB} (°C)
HFE 7000	Saturated Cycle	34	36	100	200	9	48	41
HFE 7000	Subcooled Cycle	28	29	150	100	5	75	70
HFE 7000	Endurance	30	30.5	40	260	160	50	44
HFE 7100	Saturated Cycle	61	62.5	15	585	20	64	61
HFE 7100	Subcooled Cycle	51	52	60	180	19	80	69
HFE 7100	Endurance	61	62.5	15	585	55	63	61

On the other hand, I *did* establish a test that at least strengthens my case for the feasibility of self-nucleation, even without the benefit of temperature-based data. While somewhat *ad hoc* and certainly far from ideal from a scientific standpoint, the quasi-static temperature-controlled test was in some ways a very clever attempt to circumvent the flaws of my heater/sensor system. It provided evidence that a surface treated with AIBN could be made to nucleate *well before* an untreated surface, even if neither the heat flux nor surface temperature could be measured explicitly. This alone is reasonable evidence that the superheat required for nucleation was in fact decreased.

Note the apparent slope change in Fig. 6.15 at approximately 71°C. Until this point, the temperature vs. power curve is remarkably linear. And while certainly not as pronounced as the typical “boiling plateau” associated with true boiling curves (see Fig. 3.5), it appears to indicate that *something* is happening at this point. While this slope change does not occur exactly at the temperature at which I visually observed nucleation (approximately 73°C), it is very close. Moreover, given the difficulty of pinpointing nucleation with the naked eye in such a strongly subcooled refrigerant (see Section VII.2.3), it would not surprise me if visually imperceptible nucleation first occurred on the AIBN-treated side closer to 71°C. What’s more, if one takes into account the thermal resistance of the butyrate membrane in combination with the thermistor-recorded nucleation temperature of 73°C, it seems likely that nucleation at the evaporator surface occurred *very* close to the initial thermal decomposition of AIBN, just as envisioned in Fig. 4.4. This new heat transfer mode could certainly account for the slope change. Unfortunately, I did not continue to record data after nucleate boiling was initiated on the untreated side, so I cannot verify if there was an additional slope change following that nucleation event.

Finally, there remains the crucial issue that perhaps major temperature differences were not observed simply because a large number of nucleation sites were never fully flooded and thus deactivated. Although the experiments were designed to use highly wetting refrigerants in conjunction with extreme subcooling during cyclic boiling precisely to promote site deactivation, the low superheat observed for nucleation in all tests makes such an assumption dubious. It is probable that the crude macroscale fabrication process allowed mesoscale surface roughness effects to entirely overwhelm any of the submicron scale surface features associated with location-dependent homogeneous nucleation. Nevertheless, it was still expected that the artificial creation of a gas-liquid interface within mesoscale features would promote measurably lower-temperature nucleation.

VII.2 Experimental Apparatus Shortfalls and Proposed Modifications

As discussed in Section V.5, there were many aspects of the experimental setup designed specifically to simplify data collection and analysis, to permit the use of commercially-available macroscale heating and sensing components, and to minimize the demand for extreme experimental control from one test to the next. This was done largely to minimize costs and delay time-consuming Microlab qualification, but it was also based, to some extent, on the degree of confidence I had in the self-nucleation concept. I firmly believed that phenomenological differences between the treated and untreated surfaces would be so pronounced during cyclic boiling, that the crude, merely comparative nature of the experiments would still ably demonstrate the feasibility of the concept and thus justify further doctoral research in microscale systems.

Nevertheless, based on the quasi-static temperature-controlled tests and the discussion in Section VII.1, I still believe the theory is sound and that self-nucleation *was* indeed occurring, even if the temperature data did not reflect it. At the same time, however, it would have been extremely worthwhile to consult someone with more *experimental* experience in pool boiling to help design a system to collect more demonstrative temperature-based data. Sections VII.2.1-VII.2.3 discuss some of the shortfalls of the experimental setup, as well as possible solutions to improve the results of future experiments.

VII.2.1 Heating System

As discussed in Section VI.5, the large thermal mass and high thermal conductivity of the aluminum heating block caused tremendous difficulties with controlling heat flux. While the final version of the heater was a vast improvement over the first iteration (which had nearly three times the mass), I was still constrained by the relatively large dimensions of the macroscale ceramic heater. Prescribed heat flux could not be delivered to the boiling surface quickly, and the block took an excessively long time to cool down between heating cycles. This cooling problem was further compounded by the fact that the block was not in contact with the liquid and that most of the block's surface was insulated by air. This excessive thermal isolation was, in fact, intentional—based on initial experimental concerns. It was my original intention to minimize parasitic heat losses (*i.e.* heat transfer *not* through the evaporator surface) so that I could better estimate the total heat flux that was passing through the evaporator surface and thus construct approximate boiling curves. This concern also influenced the spring-loading design, which was left “unconstrained” to minimize conduction down into the base plate. Such concerns later proved to be unnecessary when experiments were designed to be comparative, rather than explicitly quantitative. This decision toward quantitative testing was, in turn, largely based on the large time delay associated with delivering heat to the evaporator surface. Consequently, boiling curves could only be constructed on a point-by-point basis using quasi-static temperature control (see Section VI.5), which was not practical for cyclic boiling tests.

Another problem associated with the heating system was the possibility of asymmetric heat flux. Ironically, I originally chose the single-heater, split-boss design precisely to *ensure* equal heat flux to both surfaces. At the time, I was more concerned that separate heaters might exhibit non-identical heating characteristics and thus be even more difficult to control. In hindsight, asymmetric heat flux is a much bigger problem, particularly if the experimental design is based on flux-controlled boiling with temperature measurement.

Naturally, the best way to improve the heater design is to make it as small as possible, and as close to the evaporator surface as possible. In many boiling experiments, in fact, the evaporator surface *is* the heater itself. As discussed in Section II.1.3, classical pool boiling experiments used a fine nichrome wire, while most MEMS-based heat transfer studies utilize thin-film resistors. In both cases, the thermal mass of the heater is so small that the heating and cooling can be controlled almost instantaneously by the electrical power supply.

VII.2.2 Temperature Sensing System

For the most part, the thermistors themselves performed quite well. Their small size, accuracy, and rapid thermal response were certainly sufficient for this experimental study, had they only been positioned better. Unfortunately, their relative remoteness from the evaporator surface precluded even the crudest qualitative comparison-based tests. This was mainly due to the fact that they were embedded in the surface of the thermally-massive, highly-conductive aluminum heater block. Consequently, they consistently mirrored the steadily increasing and decreasing temperatures of the heater block, rather than the temperatures of the overlying boiling surface.

This thermally non-optimal design and sensor layout was primarily chosen based on the perceived need for greater freedom in surface and housing fabrication. Essentially, I faced a decision between the following two design philosophies:

- 1) Design a sensor system that is highly integrated with the evaporator surface, but by doing so, severely limit the ability to damage, alter, or replace test surfaces.
- 2) Design a less-integrated, but adaptable permanent sensor system that can easily accommodate an assortment of cheap, varied, interchangeable test surfaces and housings.

When deciding where to machine the small pockets and lead grooves to encase the thermistors, I went back and forth several times between the heater block and the underside of the housing. At that time, however, I was still unsure as to how I would fabricate the test surfaces or pattern them with AIBN. In addition, the housing design had recently gone through several changes to address refrigerant leakage, and I was pretty sure more changes would be necessary. Finally, based on the boiling curves from 3M and my initial estimates of thermal resistance, I was still quite worried about melting the housing before nucleation even occurred. In the end, I decided it was safer to sacrifice some sensing accuracy and, by doing so, leave myself as much freedom as possible for the test surfaces. And again, because I anticipated such large differences in surface temperature between low-superheat heterogeneous nucleation and traditional homogeneous nucleation, I figured that even somewhat removed from the actual evaporator surface, the thermistors would still register significant temperature variation.

I also tried to address the concern with sensor accuracy by using plastic boiling surfaces. In addition to being cheap and homogeneous (thus easily reconfigurable and disposable), their extremely low thermal conductivity would greatly reduce lateral conduction, thus “intensifying” localized hot spots and theoretically allowing easier measurement by underlying thermistors. Even more importantly, a plastic evaporator surface would also prevent conduction from the untreated side to the AIBN-treated side, which would otherwise further attenuate any measurable temperature differences. As was the case with heaters, ideal temperature sensors would consist of thin-film thermistors patterned on the actual boiling surface. Most MEMS-based temperature measurements are made this way, with the thermistors serving as both surface heaters and temperature sensors (see Fig. 7.1).




Fig. 7.1 – Schematic of typical microfabricated thin-film thermistor (from [25]).

VII.2.3 Imaging Systems and Surface Characterization

Because the majority of my original experimental design centered upon temperature-based data (rather than visual confirmation of asymmetric nucleation), I did not put significant effort into integrating high-performance imaging systems. Photographic and video data was collected with a professional-grade Canon PowerShot G6 digital camera, but the entire setup was improvised long after-the-fact and lacked ideal positioning and high-speed capabilities. Moreover, I personally have very little experience with photography and found it surprisingly difficult to capture representative photos and videos.

Compounding this difficulty was the very nature of the pool boiling phenomena. The transition from single-phase natural convection to two-phase nucleate boiling was not as clear-cut as I had anticipated. Certainly it was not as simple as “bubbles” vs. “no bubbles”. In all tests, but particularly in highly subcooled systems, the line between rapidly-rising, swirling single-phase liquid plumes and faint, rapidly-rising jets of nearly-imperceptible bubbles was extremely ill-defined. During the earliest stages of nucleation, the jets of tiny bubbles would quickly condense and dissipate nebulously back into the swirling liquid plumes, only millimeters above the evaporator surface. This often made it extremely difficult to tell if nucleation had yet occurred. For this reason, a highly magnified view of the surface itself would have been particularly helpful. Such images could have also revealed tiny embryos of released nitrogen within surface defects prior to thermally-induced nucleation. Extremely high quality magnified imaging may also have helped determine if cavities were indeed fully wetted, or if gas or vapor embryos remained in cavities after nucleate boiling had ceased.

Finally, I believe that better surface characterization may have provided considerable insight into the failure to observe site deactivation. I honestly had no estimation of the surface roughness or defect characteristics of the aerosol polymer coating. In hindsight, this was a significant weakness, but once again, I was so confident in the theory that I believed just about any surface would be suitable for self-nucleation. As a result of my zealous enthusiasm for both the simplicity and hermeticity of comparative data, my only concern was ensuring the superficial similarity of the surfaces, rather than characterizing the surfaces themselves. It probably would have been extremely beneficial to obtain SEMs or at least profilometry data of the surfaces.

VII.3 Future Work

One of the most important lessons I learned during this study is that no amount of theory or conjecture will fully anticipate all of the problems and minutiae learned from experimentation. A clever design solution to address one particular issue will often result in several new, previously unforeseen issues. Consequently, iteration becomes one of the most useful tools in developing experimental research. At this point, I would like to offer a rough description of how I might alter my existing setup to improve future results. These modifications cover the entire spectrum, from simple adjustments that could be implemented quite easily in the short term to more sweeping changes in the heating and sensing systems that would require Microlab fabrication.

As discussed in Section VII.2.2, my first priority would be to move the temperature sensors away from the heater block. Having now firmly established surface fabrication techniques, it would not be too difficult at this point to embed thermistors within the AIBN-polymer matrix in slightly larger wells. While such a design would certainly provide much more *accurate* surface temperature data, I am still not convinced this would produce the widely different surface temperatures I envisioned in my experimental design. Because of the heater block's thermal mass and control method, I think asymmetric heat flux on a nearly isothermal surface will still be the observed phenomena. On the other hand, by implementing *additional* thermistors embedded within the wells (used in conjunction with the thermistors currently embedded in the heater block, see Fig. 7.2a), much better heat flux estimations could be made, based on the observed temperature gradient across the housing membrane. A second design, shown in Fig. 7.2b, would require a slight modification to the heater boss, as well as a spring-loaded plastic plate to hold the thermistors in good thermal contact with the backside of the housing membrane, but *out of contact* with the aluminum heater block. This would give the thermistors greater sensitivity to fluctuations in surface temperature and also create a natural thermal gradient across the membrane for visual comparative observations.




Fig. 7.2 – Schematics of modified test platforms to improve temperature data.

In earlier designs, I was not confident in my ability to effectively seal a two-piece housing against refrigerant leakage, but my successful experience with E-6000 sealant has resolved those uncertainties. The next level of complexity would then involve patterning thin-film resistors on the surface of a glass or pyrex insert before traditional CNC-milling of interspaced AIBN wells (see Fig. 7.3a). Instead of the ceramic heater, these thin-film resistors would be used to heat the evaporator surface, thus permitting more sensitive temperature measurements using the current thermistor layout. An alternative design would use thin-film resistors as sensors (*i.e.* thermistors). If the metallization process has a low enough thermal budget (*i.e.* below the bulk of AIBN's thermal decomposition range), traditional machining and AIBN screen-printing could be performed *before* the thin-film resistors are patterned. This would allow the thermistors to be patterned directly overtop the AIBN wells (see Fig. 7.3b), providing ideal surface temperature data. In this case, heat could still be provided by the aluminum heater block, while the current thermistors embedded in the heater block would be used for improving heat flux estimates.

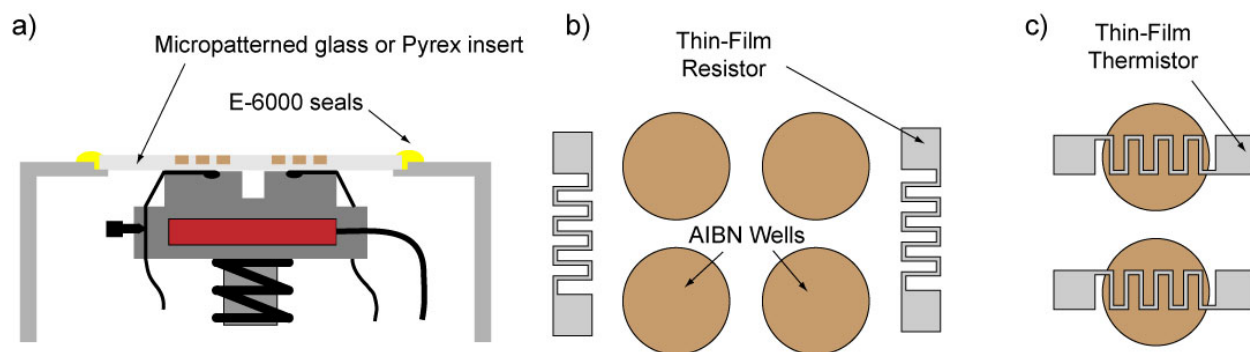


Fig. 7.3 – a) Test platform designs involving microfabricated inserts. **b)** Post-processed wells adjacent to thin-film resistors. **c)** Pre-processed wells with overlying thin-film thermistors.

Lastly, I would like to comment briefly on the future of the self-nucleating evaporator with regard to the microColumnated Loop Heat Pipe system. As emphasized in Chapter I, this study was designed primarily as an exploratory proof-of-concept for the long-term, late-phase enhancement of the μ C-LHP. Clearly, there is much work to be done before this becomes a reality. Even had this study been an unmitigated success, there are still several fundamental issues that must be resolved before a self-nucleating evaporator could be integrated into the μ C-LHP. First and foremost, it must be demonstrated that capillary pumping can provide sufficient liquid to make nucleate boiling the most desirable vaporization mode. Clearly, conditions similar to classical pool boiling will not exist inside the columnated vapor chamber. It may turn out that liquid flow rates are only sufficient to provide thin-film evaporation. However, if flow rates are large enough, a self-nucleating surface that can prevent the vapor burst commonly observed in microchannels would open up an entirely new avenue for revolutionary evaporator designs.

A second import issue is the functional lifetime of an AIBN coating. As discussed in Section IV.4, self-nucleation of deactivated sites is intended to be a rare event, more of a “safety valve” than a conventional nucleation method. While my results did not demonstrate progressively increasing nucleation temperatures (an indication of gradual nitrogen exhaustion), this was more likely due to observational problems; thus nitrogen exhaustion could have occurred without my knowledge. In any case, effective lifetime is clearly an issue of great importance and must be studied further, in terms of both the frequency of self-nucleation and the nitrogen consumption of such an event. Finally there is the issue of patterning AIBN onto an evaporator surface at the true microscale level. While encapsulation is compatible with many mainstream microfabrication polymers such as spin-on Teflon and SU-8 photoresist, patterning has thus far involved post-processed screen-printing. Consequently, the finest resolution achieved was 1 mm x 1 mm with a thickness of 75 μ m. More fabrication development is also needed before AIBN patterning can meet the demands of in-channel applications.

Appendix C: Detailed Guide for Columnated Microevaporator Fabrication

This section is designed to serve as a more concise “recipe-like” laboratory guide to supplement the fabrication description of the main text (Chapter 8). While Chapter 8 provided a general narrative of the fabrication process, it did not include many details regarding intermediate cleaning steps, specific Berkeley Microlab and Nanolab tool programs and settings, minor process variables, and other minutiae. While such details are not usually necessary to *understand* a fabrication process, they are extremely useful when trying to *replicate* one, particularly when trying to estimate fabrication *time*. As any microfabrication researcher will verify, a slight irregularity, minor process change, or even an omitted “optional” cleaning step can have dramatic effects on the final process.

That being said, this section cannot and will not cover the entire history of my fabrication experience. As mentioned at the beginning of Chapter 8, process variables for one wafer may be different from those of a second wafer, even though the final result (and thus the overall process flow) is identical. In most instances, I simply provide details *characteristic* of the final process flow. Additionally, many important procedures (*e.g.* running dummy wafers to prime photoresist lines, calculating etch rates, calibrating exposure times, *etc.*) are simply understood to be requisite to microfabrication, and such obvious steps have been omitted. To restate the example from Chapter 8, rather than documenting that I first etched for 10 minutes, *then* performed a profilometry measurement to calculate the etch rate, *and then* continued the etch to a depth of 120 microns; I simply state that I etched 120 microns. Finally, the reader is highly encouraged to continually refer back to Chapter 8 (particularly the process flow cross-sections) to maintain a firm grasp on the “big picture”; such a concept can easily be lost after the 27th Piranha clean.

Fabrication of Microevaporator Floor

Device Substrate: 6” p-type prime silicon, single-side polished, nominal thickness 625 μm

<i>Non-MOS Standard Clean</i>		
Sub-Step	Micro/Nanolab Tool	Details
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	

<i>MOS Pre-furnace Clean</i>		
Sub-Step	Micro/Nanolab Tool	Details
Piranha Clean	Sink6 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink6 QDR	4 cycles
HF Dip	Sink6 5:1 BHF (49%) Bath	1 min
Quick Dump Rinse	Sink6 QDR	4 cycles
Spin-Rinse-Dry	Sink6 SRD	

<i>LTO Deposition</i>			
Micro/Nanolab Tool	Program	Deposition Time*	Details
Tystar12	12sultoa	230 min	T = 450°C, P = 300 mT

* O₂ supply exhausted during processing; only 2-2.3 μm deposited (out of a planned 3 μm)

<i>Non-MOS Standard Clean</i>		
Sub-Step	Micro/Nanolab Tool	Details
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	
Complete Dehydration	VWR Oven	30 min, 120°C

<i>Lithography 1: Oxide Hard Mask Definition (MASK #1)</i>		
Sub-Step	Micro/Nanolab Tool	Details
HMDS	Primeoven	Prog 0, 1 min prime, 90°C
Spin-on	SVGCoat1	Spin Prog #2: 1.3μm OCG 825 (G-Line) PR 30 sec, 5000 rpm
Soft Bake	SVGCoat1	Bake Prog #1: 60 sec @ 90°C
Exposure	KSAaligner	Hard Contact Mode, 40μm gap, 8 sec exposure G-Line Lamp Intensity: 22.4 mW/cm ²
Post-Exposure Bake	SVGDev6	PEB Prog #9: No Bake Dev Prog #9: No Develop Hard Bake Prog #1: 60 sec @ 120°C
Develop	SVGDev6	PEB Prog #9: No Bake Dev Prog #2: OCG 934 2:1 (G-Line), 2 × 30 sec, 20°C Hard Bake Prog: #9: No Bake
Hard Bake	UVBake	Prog S: 3 sec FLASH @ 120°C, 10 sec OFF @ 240°C, 10 sec LOW @ 240°C, 120 sec HIGH @ 240°C

<i>Oxide Hard Mask Patterning</i>			
Micro/Nanolab Tool	Program	Time	Details
CenturaMXP	MXP-OXIDE-VAR - CF ₄ (15 sccm) - CHF ₃ (45 sccm) - Ar (150 sccm)	345 sec	B = 30 gauss (sine) RF Power = 700W P = 200 mT

<i>Post-Etch Descum</i>			
Micro/Nanolab Tool	Gas	Time	Details
Technics-C	O ₂ (51.1 sccm)	60 sec	RF Power = 50W, P = 270 mT

<i>Photomask Strip and Clean</i>		
Sub-Step	Micro/Nanolab Tool	Details
Strip PR	Sink5 Long Soak	PRS-3000, 80°C, 3 hours
Quick Dump Rinse	Sink5 QDR	4 cycles
Spin-Rinse-Dry	Sink5 SRD	
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	
Complete Dehydration	VWR Oven	30 min, 120°C

<i>Lithography 2: Thick Photoresist Soft Mask Patterning (MASK #2)</i>		
Sub-Step	Micro/Nanolab Tool	Details
HMDS	Primeoven	Prog 0, 1 min prime, 90°C
Spin-on	SVGCoat1	Spin Prog #4: 2.0µm OCG 825 (G-Line) PR 30 sec, 2200 rpm
Soft Bake	SVGCoat1	Bake Prog #2: 90 sec @ 90°C
REPEAT Spin-On and Soft Bake × 4 (5 coats total)		
Exposure	KSAligner	Hard Contact Mode, 30µm gap, 55 sec exposure G-Line Lamp Intensity: 22.8 mW/cm ²
Develop	SVGDev6	PEB Prog #9: No Bake Dev Prog #2: OCG 934 2:1 (G-Line), 2 × 30 sec, 20°C Hard Bake Prog: #9: No Bake
REPEAT Develop (2 washes total)		
Hard Bake	VWR	5 hours @ 120°C

<i>Backside Oxide Removal</i>			
Micro/Nanolab Tool	Etchant	Time	Details
Msink5	5:1 BHF (49%)	35 min	“Floating Etch” (see Chapter 8, Fig. 8.6)

<i>DRIE Partial Through-Hole Etch</i>				
Micro/Nanolab Tool	Program	Time	Chuck Temperature	RF Coil Power
STS	HEXA100	175 min*	45°C	600W
Process Variable	Etch Cycle		Passivation Cycle	
Gases	SF ₆ (130 sccm), O ₂ (13 sccm)		C ₄ F ₈ (85 sccm)	
Cycle Time	9.0 sec		7.0 sec	
RF Platen Power	12W		0W	
Over Run	0.5 sec		0.5 sec	

* Etching done in 5 steps with inspection after each step: 100 min, 60 min, 5 min, 5 min, 5 min

Thick Photoresist Soft Mask Strip and Clean

Sub-Step	Micro/Nanolab Tool	Details
Strip PR	Sink5 Long Soak	PRS-3000, 80°C, 8 hours
Quick Dump Rinse	Sink5 QDR	4 cycles
Spin-Rinse-Dry	Sink5 SRD	
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	

Handle Wafer Substrate: 6" p-type prime silicon, single-side polished, nominal thickness 625 μm

Handle Wafer Non-MOS Standard Clean

Sub-Step	Micro/Nanolab Tool	Details
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	

Handle Wafer MOS Pre-furnace Clean

Sub-Step	Micro/Nanolab Tool	Details
Piranha Clean	Sink6 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink6 QDR	4 cycles
HF Dip	Sink6 5:1 BHF (49%) Bath	1 min
Quick Dump Rinse	Sink6 QDR	4 cycles
Spin-Rinse-Dry	Sink6 SRD	

Handle Wafer LTO Deposition

Micro/Nanolab Tool	Program	Deposition Time	Details
Tystar12	12sulto	150 min	T = 450°C, P = 300 mT

Reversible Handle Wafer Bond

Sub-Step	Micro/Nanolab Tool	Details
Spin-on (Handle)	SVGCoat1	Spin Prog #4: 2.0 μm OCG 825 (G-Line) PR 30 sec, 2200 rpm
Soft Bake	SVGCoat1	Bake Prog #9: No Bake
Bond	N/A	Press Together (see Chapter 8, Fig. 8.11, Step 12)
Hard Bake	Primeoven	Prog 0, 90°C, orient wafers horizontally
Pre-STs Pump Down	Technics-C	No gases, P = 30 mT, 15 min

<i>DRIE Hole Punch-Through and Channel Etch</i>				
Micro/Nanolab Tool	Program	Time	Chuck Temperature	RF Coil Power
STS	AARON1B	60 min	45°C	600W
Process Variable	Etch Cycle		Passivation Cycle	
Gases	SF ₆ (130 sccm), O ₂ (13 sccm)		C ₄ F ₈ (85 sccm)	
Cycle Time	10.0 sec		8.0 sec	
RF Platen Power	12W		0W	
Over Run	0.5 sec		0.5 sec	

<i>Handle Wafer De-Bond</i>		
Sub-Step	Micro/Nanolab Tool	Details
Dissolve PR	Sink5 Long Soak	PRS-3000, 80°C, 8 hours
De-bond	N/A	Use razor blade to slide acetone-soaked TechniCloth between wafers

<i>Post-De-Bond Clean</i>		
Sub-Step	Micro/Nanolab Tool	Details
Quick Dump Rinse	Sink5 QDR	4 cycles
Spin-Rinse-Dry	Sink5 SRD	
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles

<i>Oxide Hard Mask Removal</i>		
Sub-Step	Micro/Nanolab Tool	Details
Oxide Etch	Sink8 5:1 BHF (49%) Bath	20 min
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	

<i>Wafer Dice</i>				
Micro/Nanolab Tool	Blade	Spindle Speed	Feed Rate	Blade-to-Chuck
Wafersaw	K3T20L45	30,000 rpm	1 mm/s	150 µm

Fabrication of Microevaporator Ceiling

Device Substrate: 6" p-type prime silicon, single-side polished, nominal thickness 625 μm

<i>Non-MOS Standard Clean</i>		
Sub-Step	Micro/Nanolab Tool	Details
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles
HF Dip	Sink8 5:1 BHF (49%) Bath	1 min
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	
Complete Dehydration	VWR Oven	30 min, 120°C

<i>Lithography 3: Photoresist Mask Definition (MASK #3)</i>		
Sub-Step	Micro/Nanolab Tool	Details
HMDS	Primeoven	Prog 0, 1 min prime, 90°C
Spin-on	SVGCoat1	Spin Prog #1: 1.1 μm OCG OiR 10i (I-Line) PR 30 sec, 4100 rpm
Soft Bake	SVGCoat1	Bake Prog #1: 60 sec @ 90°C
Exposure	KSAaligner	Vacuum Contact Mode, 40 μm gap, 8 sec exposure I-Line Lamp Intensity: 9.2-9.6 mW/cm ²
Post-Exposure Bake	SVGDev6	PEB Prog #9: No Bake Dev Prog #9: No Develop Hard Bake Prog #1: 60 sec @ 120°C
Develop	SVGDev6	PEB Prog #9: No Bake Dev Prog #3: OPD 4262 (I-Line), 60 sec, 20°C Hard Bake Prog: #9: No Bake
Hard Bake	UVBake	Prog J: 0 sec FLASH @ 110°C, 10 sec OFF @ 230°C, 10 sec LOW @ 230°C, 70 sec HIGH @ 230°C

<i>DRIE Channel Etch</i>				
Micro/Nanolab Tool	Program	Time*	Chuck Temperature	RF Coil Power
STS	HEXA100	10:40	45°C	600W
Process Variable	Etch Cycle		Passivation Cycle	
Gases	SF ₆ (130 sccm), O ₂ (13 sccm)		C ₄ F ₈ (85 sccm)	
Cycle Time	9.0 sec		7.0 sec	
RF Platen Power	12W		0W	
Over Run	0.5 sec		0.5 sec	

* Etching done in 2 steps of 5:20 with ASIQ step height measurement between

<i>Photoresist Mask Strip and Clean</i>		
Sub-Step	Micro/Nanolab Tool	Details
Strip PR	Sink5 Long Soak	PRS-3000, 80°C, 3 hours
Quick Dump Rinse	Sink5 QDR	4 cycles
Spin-Rinse-Dry	Sink5 SRD	
Piranha Clean	Sink8 Piranha Bath	120°C, 10 minutes
Quick Dump Rinse	Sink8 QDR	4 cycles
Spin-Rinse-Dry	Sink8 SRD	

<i>Wafer Dice</i>				
Micro/Nanolab Tool	Blade	Spindle Speed	Feed Rate	Blade-to-Chuck
Wafersaw	K3T20L45	30,000 rpm	1 mm/s	150 µm

Appendix D: Sample Calculation of Data Conversion

In this subsection, I demonstrate the data conversion technique outlined in the previous subsection with a complete sample calculation using simulated data. For the sake of clarity, I have simplified the geometry of the clamps to be simple rectangular prisms, but the calculation method is the same, regardless of the number and precise areas of the various exposed surfaces. Fig. D.1 below summarizes the nomenclature and specifies the origin and values of the raw data.




Fig. D.1 – Summary of nomenclature and raw data values for sample calculation.

1. Calculate the approximate temperature gradient between the two lines at the center of evaporator.

$$\frac{dT}{dx} \approx \frac{\Delta T}{\Delta x} = \frac{T_{\text{Line } 1} - T_{\text{Line } 2}}{\Delta x} = \frac{110.1 - 108.9}{0.003} = 400 \text{ K/m}$$

2. Calculate the film temperature for each of the clamps as defined by Eq. 8.11 (assuming constant surface temperature equal to the mean temperature within each Box).

$$T_{f,HC} = \frac{T_{\text{Box } 1} + T_{\text{ambient}}}{2} = \frac{95.8 + 27.0}{2} = 61.4^\circ \text{C}$$

$$T_{f,DC} = \frac{T_{\text{Box } 2} + T_{\text{ambient}}}{2} = \frac{64.3 + 27.0}{2} = 45.7^\circ \text{C}$$

3. Evaluate the relevant thermophysical properties of the surrounding air at each clamp's film temperature.

Surface	T_f (°C)	β (K ⁻¹)	ν (m ² /s)	α (m ² /s)	k (W/m·K)	Pr
HC	61.4	0.00299	19.37×10^{-6}	27.61×10^{-6}	28.86×10^{-3}	0.701
DC	45.7	0.00314	17.78×10^{-6}	25.28×10^{-6}	27.69×10^{-3}	0.703

4. Calculate the characteristic length of each exposed surface. For vertical surfaces (*i.e.* front, back, and sides), the correlation (Eq. 8.14) uses a Rayleigh number where L is defined as the length of the vertical surface. For horizontal surfaces (*i.e.* top and bottom), the correlations (Eqs. 8.12 – 8.13) use a Rayleigh number where L is defined as the surface area divided by the perimeter.

$$L_{HC,top\&bottom} = \frac{A_{HC,top}}{P_{HC,top}} = \frac{L4 \times L6}{2(L4 + L6)} = \frac{(0.017)(0.032)}{2(0.017 + 0.032)} = 0.00555 \text{ m}$$

$$L_{HC,front,back,sides} = L5 = 0.014 \text{ m}$$

$$L_{DC,top\&bottom} = \frac{A_{DC,top}}{P_{DC,top}} = \frac{L1 \times L3}{2(L1 + L3)} = \frac{(0.014)(0.0305)}{2(0.014 + 0.0305)} = 0.00480 \text{ m}$$

$$L_{DC,front,back,sides} = L2 = 0.016 \text{ m}$$

5. Now calculate the Rayleigh number for each exposed surface, using the thermophysical property values found in Step 3.

$$Ra_{L,HC} = \frac{g \beta (T_{HC} - T_{ambient}) L^3}{\nu \alpha} = \frac{(9.8)(0.00299)(95.8 - 27.0)L^3}{(19.37 \times 10^{-6})(27.61 \times 10^{-6})} = 3.773 \times 10^9 L^3$$

$$Ra_{L,HC,top\&bottom} = (3.773 \times 10^9)(0.00555)^3 = 645$$

$$Ra_{L,HC,front,back,sides} = (3.773 \times 10^9)(0.014)^3 = 10352$$

$$Ra_{L,DC} = \frac{g \beta (T_{DC} - T_{ambient}) L^3}{\nu \alpha} = \frac{(9.8)(0.00314)(64.3 - 27.0)L^3}{(17.78 \times 10^{-6})(25.28 \times 10^{-6})} = 2.556 \times 10^9 L^3$$

$$Ra_{L,DC,top\&bottom} = (2.556 \times 10^9)(0.00480)^3 = 282$$

$$Ra_{L,DC,front,back,sides} = (2.556 \times 10^9)(0.016)^3 = 10458$$

6. Now use the proper correlation from Eqs. 8.12 – 8.14 (taking into account the orientation of the surface) to calculate the average Nusselt number for each surface.

$$\overline{Nu}_{L,HC,top} = 0.54 (645)^{1/4} = 2.72$$

$$\overline{Nu}_{L,HC,bottom} = 0.27 (645)^{1/4} = 1.36$$

$$\overline{Nu}_{L,HC,front,back,sides} = 0.68 + \frac{(0.67)(10352)^{1/4}}{[1 + (0.492/0.701)^{9/16}]^{4/9}} = 5.86$$

$$\overline{Nu}_{L,DC,top} = 0.54 (282)^{1/4} = 2.21$$

$$\overline{Nu}_{L,DC,bottom} = 0.27 (282)^{1/4} = 1.11$$

$$\overline{Nu}_{L,DC,front,back,sides} = 0.68 + \frac{(0.67)(10458)^{1/4}}{[1 + (0.492/0.703)^{9/16}]^{4/9}} = 5.87$$

7. For each surface and corresponding average Nusselt number, extract the average heat transfer coefficient h_{AVG} (noting that k_{fluid} is again evaluated at the film temperature).

$$\overline{Nu}_L = \frac{h_{AVG} \cdot L}{k_{fluid}} \Rightarrow h_{AVG} = \frac{\overline{Nu}_L \cdot k_{fluid}}{L}$$

$$h_{AVG,HC,top} = \frac{(2.72)(28.86 \times 10^{-3})}{0.00555} = 14.1 \text{ W/m}^2 \cdot \text{K}$$

$$h_{AVG,HC,bottom} = \frac{(1.36)(28.86 \times 10^{-3})}{0.00555} = 7.07 \text{ W/m}^2 \cdot \text{K}$$

$$h_{AVG,HC,front,back,sides} = \frac{(5.86)(28.86 \times 10^{-3})}{0.014} = 12.1 \text{ W/m}^2 \cdot \text{K}$$

$$h_{AVG,DC,top} = \frac{(2.21)(27.69 \times 10^{-3})}{0.00480} = 12.8 \text{ W/m}^2 \cdot \text{K}$$

$$h_{AVG,DC,bottom} = \frac{(1.11)(27.69 \times 10^{-3})}{0.00480} = 6.39 \text{ W/m}^2 \cdot \text{K}$$

$$h_{AVG,DC,front,back,sides} = \frac{(5.87)(27.69 \times 10^{-3})}{0.016} = 10.2 \text{ W/m}^2 \cdot \text{K}$$

8. Now calculate the heat lost to natural convection from each surface. Multiply the average heat transfer coefficient times the corresponding surface area times the temperature difference between the surface and the ambient.

$$Q = h A_{surface} (T_{surface} - T_{ambient})$$

$$Q_{HC,top} = (14.1)(0.017)(0.032)(95.8 - 27.0) = 0.530 \text{ W}$$

$$Q_{HC,bottom} = (7.07)(0.017)(0.032)(95.8 - 27.0) = 0.265$$

$$Q_{HC,front\&back} = 2 \times (12.1 \text{ W/m}^2 \cdot \text{K})(0.017 \text{ m})(0.014 \text{ m})(95.8 - 27.0) = 0.395 \text{ W}$$

$$Q_{HC,sides} = 2 \times (12.1 \text{ W/m}^2 \cdot \text{K})(0.014 \text{ m})(0.032 \text{ m})(95.8 - 27.0) = 0.744 \text{ W}$$

$$Q_{HC,TOTAL} = 0.528 + 0.265 + 0.395 + 0.744 = 1.93 \text{ W}$$

$$Q_{DC,top} = (12.8)(0.014)(0.0305)(64.3 - 27.0) = 0.203 \text{ W}$$

$$Q_{DC,bottom} = (6.39)(0.014)(0.0305)(64.3 - 27.0) = 0.102 \text{ W}$$

$$Q_{DC,front\&back} = 2 \times (10.2 \text{ W/m}^2 \cdot \text{K})(0.014 \text{ m})(0.016 \text{ m})(64.3 - 27.0) = 0.170 \text{ W}$$

$$Q_{DC,sides} = 2 \times (10.2 \text{ W/m}^2 \cdot \text{K})(0.016 \text{ m})(0.0305 \text{ m})(64.3 - 27.0) = 0.370 \text{ W}$$

$$Q_{DC,TOTAL} = 0.203 + 0.102 + 0.170 + 0.370 = 0.845 \text{ W}$$

9. Now add the convective losses from each clamp to obtain an estimate for the total heat lost to convection.

$$Q_{lost\ to\ convection} = 1.93 \text{ W} + 0.845 \text{ W} = 2.78 \text{ W}$$

10. The total heat flowing through the device at steady state is the heat dissipated by the heater (voltage \times current) *minus* the heat lost to convection.

$$Q_{conducted\ through\ device} = Q_{dissipated\ by\ heater} - Q_{lost\ to\ convection}$$

$$Q_{conducted\ through\ device} = (19.7 \text{ V})(0.274 \text{ A}) - 2.78 \text{ W} = 2.62 \text{ W}$$

11. Following Eq. 8.6, use the approximate temperature gradient (from Step 1), the total conducted heat (from Step 10), and the cross sectional area of the device to calculate the effective thermal conductivity.

$$k_{EFF} = \frac{Q}{A_{cross}} \frac{\Delta x}{\Delta T} = \frac{2.62 \text{ W}}{(0.014 \text{ m})(0.0012 \text{ m})} \frac{1}{400 \text{ K/m}} = 390 \text{ W/m} \cdot \text{K}$$

Alternate Method: Spherical Clamp Approximation

As discussed in the previous subsection, I also performed a second convective loss calculations that approximates the clamps as spheres, in order to utilize a correlation without lower limits on applicable Rayleigh number. Steps 1-3 in the above process remain unchanged.

- 4A. The characteristic length for the sphere correlation (Eq. 8.15) is the diameter. Clearly there is no obvious choice for diameter in a rectangular prism, so I have used the cube root of the volume, $(L \times W \times H)^{1/3}$.

$$D_{HC} = \sqrt[3]{L4 \times L5 \times L6} = \sqrt[3]{(0.017)(0.014)(0.032)} = 0.0197 \text{ m}$$

$$D_{DC} = \sqrt[3]{L1 \times L2 \times L3} = \sqrt[3]{(0.014)(0.016)(0.0305)} = 0.0190 \text{ m}$$

5A. Now calculate the Rayleigh number for each “spherical” clamp based on this diameter.

$$Ra_{D,HC} = \frac{g \beta (T_{HC} - T_{ambient}) D^3}{\nu \alpha} = \frac{(9.8)(0.00299)(95.8 - 27.0)(0.0197)^3}{(19.37 \times 10^{-6})(27.61 \times 10^{-6})} = 28732$$

$$Ra_{D,DC} = \frac{g \beta (T_{DC} - T_{ambient}) D^3}{\nu \alpha} = \frac{(9.8)(0.00314)(64.3 - 27.0)(0.0190)^3}{(17.78 \times 10^{-6})(25.28 \times 10^{-6})} = 17443$$

6A. Use Eq. 8.15 to find the average Nusselt numbers.

$$\overline{Nu}_{D,HC} = 2 + \frac{0.589 (28732)^{1/4}}{[1 + (0.469/0.701)^{9/16}]^{4/9}} = 7.91$$

$$\overline{Nu}_{D,DC} = 2 + \frac{0.589 (17443)^{1/4}}{[1 + (0.469/0.703)^{9/16}]^{4/9}} = 7.22$$

7A. Now extract the average heat transfer coefficients.

$$h_{AVG} = \frac{\overline{Nu}_L \cdot k_{fluid}}{L}$$

$$h_{AVG,HC} = \frac{(7.91)(28.86 \times 10^{-3})}{0.0197} = 11.6 \text{ W/m}^2 \cdot \text{K}$$

$$h_{AVG,DC} = \frac{(7.22)(27.69 \times 10^{-3})}{0.0190} = 10.5 \text{ W/m}^2 \cdot \text{K}$$

8A. Now find the convective losses, using the entire surface area of each prism.

$$A_{HC} = 2 \times [(0.017)(0.032) + (0.014)(0.017) + (0.014)(0.032)] = 0.00246 \text{ m}^2$$

$$Q_{HC} = (11.6 \text{ W/m}^2 \cdot \text{K})(0.00246 \text{ m}^2)(95.8 - 27.0) = 1.96 \text{ W}$$

$$A_{DC} = 2 \times [(0.017)(0.032) + (0.014)(0.017) + (0.014)(0.032)] = 0.00228 \text{ m}^2$$

$$Q_{DC} = (10.5 \text{ W/m}^2 \cdot \text{K})(0.00228 \text{ m}^2)(64.3 - 27.0) = 0.895 \text{ W}$$

9A. Add the convective losses from both clamps to find the total heat loss.

$$Q_{lost\ to\ convection} = 1.96\ W + 0.895\ W = 2.85\ W$$

10A. Now subtract this from the dissipated power to find the heat flow through the device.

$$Q_{conducted\ through\ device} = (19.7\ V)(0.274\ A) - 2.85\ W = 2.55\ W$$

11A. Finally, calculate the effective thermal conductivity as in Step 11 above.

$$k_{EFF} = \frac{Q}{A_{cross}} \frac{\Delta x}{\Delta T} = \frac{2.55\ W}{(0.014\ m)(0.0012\ m)} \frac{1}{400\ K/m} = 379\ W/m \cdot K$$

Note that this value, using a very rough spherical clamp approximation, differs by less than 3% from the original method. Moreover, this exceptional agreement was not due strictly to chance (*i.e.* overestimated losses on one clamp cancelling out *underestimated* losses on the other). The sphere-approximated losses for each clamp *individually* were extremely close to the original values (1.93 vs. 1.96 W and 0.844 vs. 0.895 W, respectively). This gives me some level of confidence in the use of the horizontal plate correlations, even though the Rayleigh numbers are below the suggested limits. This also suggests that for such small surfaces and temperature differences (and thus correspondingly low Rayleigh numbers), geometry is not so critical. This also makes some intuitive sense; as the dimensions of the clamps shrink, they appear as mere point sources of natural convection (or *plumes* in the engineering parlance).

Appendix E: Device Identification Key

Tables E.1 and E.2 enumerate the device identification codes used throughout Chapter 8. In device designation, the Floor ID precedes the Ceiling ID. Thus device B2C3 is composed of Floor B2 (column width = 500 μm , column pitch = 2500 μm) and Ceiling C3 (orthogonal 8 μm channels with 4 μm spacing). Device locations on the wafer are illustrated in Figs. E.1 and E.2.

Table E.1 – Device Floor Identification Codes

Floor ID	Column Width (μm)	Column Pitch (μm)	Floor ID	Column Width (μm)	Column Pitch (μm)
A1	200	2500	C1	900	5000
A2	500	1250	C2	500	2500
A3	900	5000	C3	200	2500
A4	300	2500	C4	900	1250
A5	300	1250	C5	300	5000
A6	500	2500	C6	200	5000
A7	300	1250	C7	300	2500
A8	200	5000	C8	500	5000
B1	900	2500	D1	500	1250
B2	500	2500	D2	300	2500
B3	900	5000	D3	200	1250
B4	900	1250	D4	300	1230
B5	200	1250	D5	900	2500
B6	300	5000	D6	200	5000
B7	200	1250	D7	900	2500
B8	500	5000	D8	500	1250
			E1	900	1250




Fig. E.1 – Device Floor Wafer Map.

Table. E.2 – Device Ceiling Identification Codes

Ceiling ID	Scheme	Channel Width (μm)	Channel Spacing (μm)
A1	Parallel	16	8
A2	Orthogonal	8	8
A3	Parallel	16	16
A4	Parallel	8	8
A5	Parallel	2	2
A6	Orthogonal	2	2
B1	Orthogonal	16	8
B2	Parallel	2	4
B3	Orthogonal	2	4
B4	Orthogonal	16	16
B5	Parallel	4	4
B6	Orthogonal	4	4
C1	Oblique	2	2
C2	Parallel	8	4
C3	Orthogonal	8	4
C4	Oblique	4	4
C5	Parallel	8	8
C6	Orthogonal	8	8
D1	Oblique	8	8
D2	Parallel	16	4
D3	Orthogonal	16	4
D4	Oblique	16	16




Fig. E.2 – Device Ceiling Wafer Map.

Appendix F: Heat Transfer Analysis MATLAB Code

Thermophysical Property Calculation Functions

```
function [Tfilm] = Tfilm(Tsurf,Tamb)
Tfilm = 0.5*(Tsurf+Tamb);
```

```
function [alphaTH] = alphaTH(T)
Tkelvin = T + 273.15;
alphaTH = (22.5+(Tkelvin-300)*(29.9-22.5)/(350-300))*10^(-6);
```

```
function [betaTH] = betaTH(T)
Tkelvin = T + 273.15;
betaTH = 1/Tkelvin;
```

```
function [Kfluid] = Kfluid(T)
Tkelvin = T + 273.15;
Kfluid = (26.3+(Tkelvin-300)*(30-26.3)/(350-300))*10^(-3);
```

```
function [v] = viscosity(T)
Tkelvin = T + 273.15;
v = (15.89+(Tkelvin-300)*(20.92-15.89)/(350-300))*10^(-6);
```

```
function [Pr] = Pr(T)
Pr = viscosity(T)/alphaTH(T);
```

Rayleigh Number Calculator

```
function [Ra] = Ra(Tsurf,Tamb,Lchar)
Tf = Tfilm(Tsurf,Tamb);
Ra = (9.81*betaTH(Tf)*(Tsurf-Tamb)*Lchar^3)/(viscosity(Tf)*alphaTH(Tf));
```

Nusselt Number Correlations

```
function [Nu] = NuTOP(Ra)
Nu = 0.54*Ra^0.25;
```

```
function [Nu] = NuBOT(Ra)
Nu = 0.27*Ra^0.25;
```

```
function [Nu] = NuLAT(Ra,Pr)
Nu = 0.68+(0.67)*Ra^0.25/(1+(0.492/Pr)^(9/16))^(4/9);
```

Convection Coefficient Extrator

```
function [hCONV] = hCONV(Nu,kFL,L);
hCONV = Nu*kFL/L;
```

General Isothermal Prism Convective Loss Calculator

```
function [qTOP,qBOT,qLAT,qTOT] = prismCONV_isot(L,W,H,Tsurf,Tamb)
Tf = Tfilm(Tsurf,Tamb);
```

```
%Top and Bottom Surfaces
```

```
A_TB = L*W;
P_TB = 2*(L+W);
Lchar_TB = A_TB/P_TB;
```

```
%Top Surface
```

```
Ra1 = Ra(Tsurf,Tamb,Lchar_TB);
Nu1 = NuTOP(Ra1);
h1 = hCONV(Nu1,Kfluid(Tf),Lchar_TB);
qTOP = h1*A_TB*(Tsurf-Tamb);
```

```
%Bottom Surface
```

```
Ra2 = Ra(Tsurf,Tamb,Lchar_TB);
Nu2 = NuBOT(Ra2);
h2 = hCONV(Nu2,Kfluid(Tf),Lchar_TB);
qBOT = h2*A_TB*(Tsurf-Tamb);
```

```
%Lateral Surfaces
```

```
Lchar3 = H;
A3 = 2*(L*H+W*H);
Ra3 = Ra(Tsurf,Tamb,Lchar3);
Nu3 = NuLAT(Ra3,Pr(Tf));
h3 = hCONV(Nu3,Kfluid(Tf),Lchar3);
qLAT = h3*A3*(Tsurf-Tamb);
```

```
qTOT = qTOP+qBOT+qLAT;
```

Corrective Clamp Surface Temperature Correlations

```
HCtopC=[56.5 60.8 67.8 73.5 78.4 84 90.5];
HCtopE=[49.1 52 56.9 60.6 64.6 69.1 73.1];
HCbotC=[56.8 59.6 65.3 70.6 75.8 81.2 86.3];
HCbotE=[49.2 51.3 54.7 58.5 60.1 65.6 68.2];
HCfrontTop=[52.6 57 61.4 65.5 72 75.2 79.3];
HCfrontBot=[49.8 53.8 57.7 61.5 67.2 70.1 73.9];
```

```
DCTopC=[45.9 48.4 53.6 57.3 60.2 63.7 66.8];
DCTopE=[40.4 42.7 47.1 50.1 52.4 55.7 58.8];
DCbotC=[47.6 50 54.9 58.6 62.4 66.1 69.9];
DCbotE=[40.9 43.5 46.6 49.5 49.9 51.6 55.9];
DCfrontTop=[42.3 44.8 47.1 49.4 52.8 54.4 56.4];
DCfrontBot=[41.8 44.2 46.3 48.7 51.8 53.4 55.4];
```

```
HCtopE_coeff = polyfit(HCtopC,HCtopE,1);
HCbotC_coeff = polyfit(HCtopC,HCbotC,1);
HCbotE_coeff = polyfit(HCtopC,HCbotE,1);
HCfrontTop_coeff = polyfit(HCtopC,HCfrontTop,1);
HCfrontBot_coeff = polyfit(HCtopC,HCfrontBot,1);
```

```
DCTopE_coeff = polyfit(DCTopC,DCTopE,1);
DCbotC_coeff = polyfit(DCTopC,DCbotC,1);
DCbotE_coeff = polyfit(DCTopC,DCbotE,1);
DCfrontTop_coeff = polyfit(DCTopC,DCfrontTop,1);
DCfrontBot_coeff = polyfit(DCTopC,DCfrontBot,1);
```

General Correlated Weighted Prism Convective Loss Calculator

```
function [qTOP,qBOT,qLAT,qTOT] = prismCONV_weighted(L,W,H,Tsurf,Tamb,Cweight)

%Top and Bottom Surface Geometry
A_TB = L*W;
P_TB = 2*(L+W);
Lchar_TB = A_TB/P_TB;

%Top Surface Losses
Tcent1 = Tsurf;
Tedge1 = 0.7159*Tsurf+8.4628;
Tfcent1 = Tfilm(Tcent1,Tamb);
Tfedge1 = Tfilm(Tedge1,Tamb);
Ra1C = Ra(Tcent1,Tamb,Lchar_TB);
Ra1E = Ra(Tedge1,Tamb,Lchar_TB);
Nu1C = NuTOP(Ra1C);
Nu1E = NuTOP(Ra1E);
h1C = hCONV(Nu1C,Kfluid(Tfcent1),Lchar_TB);
h1E = hCONV(Nu1E,Kfluid(Tfedge1),Lchar_TB);
qTOP = A_TB*(h1C*Cweight*(Tcent1-Tamb)+h1E*(1-Cweight)*(Tedge1-Tamb));
```

```

%Bottom Surface Losses
Tcent2 = 0.8939*Tsurf+5.4790;
Tedge2 = 0.5722*Tsurf+16.4191;
Tfcent2 = Tfilm(Tcent2,Tamb);
Tfedge2 = Tfilm(Tedge2,Tamb);
Ra2C = Ra(Tcent2,Tamb,Lchar_TB);
Ra2E = Ra(Tedge2,Tamb,Lchar_TB);
Nu2C = NuBOT(Ra2C);
Nu2E = NuBOT(Ra2E);
h2C = hCONV(Nu2C,Kfluid(Tfcent2),Lchar_TB);
h2E = hCONV(Nu2E,Kfluid(Tfedge2),Lchar_TB);
qBOT = A_TB*(h2C*Cweight*(Tcent2-Tamb)+h2E*(1-Cweight)*(Tedge2-Tamb));

%Lateral Surfaces Losses
A_LAT = 2*(L*H+W*H);
TlatTOP = 0.7967*Tsurf+7.9244;
TlatBOT = 0.7172*Tsurf+9.5940;
TflatTOP = Tfilm(TlatTOP,Tamb);
TflatBOT = Tfilm(TlatBOT,Tamb);
Lchar3 = H;
A3 = 2*(L*H+W*H);
Ra3TOP = Ra(TlatTOP,Tamb,Lchar3);
Ra3BOT = Ra(TlatBOT,Tamb,Lchar3);
Nu3TOP = NuLAT(Ra3TOP,Pr(TflatTOP));
Nu3BOT = NuLAT(Ra3BOT,Pr(TflatBOT));
h3TOP = hCONV(Nu3TOP,Kfluid(TflatTOP),Lchar3);
h3BOT = hCONV(Nu3BOT,Kfluid(TflatBOT),Lchar3);
qLAT = A_LAT*(h3TOP*0.5*(TlatTOP-Tamb)+h3BOT*0.5*(TlatBOT-Tamb));

qTOT = qTOP+qBOT+qLAT;

```

Correlated Weighted Device Clamp Convective Loss Calculator

```

function [qTOP,qBOT,qLAT,qTOT] = qDCweighted(Tsurf,Tamb,Cweight)

L = 0.014;
W = 0.0305;
H = 0.016;

%Top and Bottom Surface Geometry
A_TB = L*W;
P_TB = 2*(L+W);
Lchar_TB = A_TB/P_TB;

%Top Surface Losses
Tcent1 = Tsurf;
Tedge1 = 0.8647*Tsurf+0.6944;
Tfcent1 = Tfilm(Tcent1,Tamb);
Tfedge1 = Tfilm(Tedge1,Tamb);
Ra1C = Ra(Tcent1,Tamb,Lchar_TB);
Ra1E = Ra(Tedge1,Tamb,Lchar_TB);
Nu1C = NuTOP(Ra1C);

```

```

Nu1E = NuTOP(Ra1E);
h1C = hCONV(Nu1C,Kfluid(Tfcent1),Lchar_TB);
h1E = hCONV(Nu1E,Kfluid(Tfedge1),Lchar_TB);
qTOP = A_TB*(h1C*Cweight*(Tcent1-Tamb)+h1E*(1-Cweight)*(Tedge1-Tamb));

%Bottom Surface Losses
Tcent2 = 1.0638*Tsurf-1.6628;
Tedge2 = 0.6446*Tsurf+11.8120;
Tfcent2 = Tfilm(Tcent2,Tamb);
Tfedge2 = Tfilm(Tedge2,Tamb);
Ra2C = Ra(Tcent2,Tamb,Lchar_TB);
Ra2E = Ra(Tedge2,Tamb,Lchar_TB);
Nu2C = NuBOT(Ra2C);
Nu2E = NuBOT(Ra2E);
h2C = hCONV(Nu2C,Kfluid(Tfcent2),Lchar_TB);
h2E = hCONV(Nu2E,Kfluid(Tfedge2),Lchar_TB);
qBOT = A_TB*(h2C*Cweight*(Tcent2-Tamb)+h2E*(1-Cweight)*(Tedge2-Tamb));

%Lateral Surfaces Losses
A_3 = 2*(L*H+W*H);
T3TOP = 0.6687*Tsurf+11.7830;
T3BOT = 0.6437*Tsurf+12.3960;
Tf3TOP = Tfilm(T3TOP,Tamb);
Tf3BOT = Tfilm(T3BOT,Tamb);
Lchar3 = H;
Ra3TOP = Ra(T3TOP,Tamb,Lchar3);
Ra3BOT = Ra(T3BOT,Tamb,Lchar3);
Nu3TOP = NuLAT(Ra3TOP,Pr(Tf3TOP));
Nu3BOT = NuLAT(Ra3BOT,Pr(Tf3BOT));
h3TOP = hCONV(Nu3TOP,Kfluid(Tf3TOP),Lchar3);
h3BOT = hCONV(Nu3BOT,Kfluid(Tf3BOT),Lchar3);
qLAT = A_3*(h3TOP*0.5*(T3TOP-Tamb)+h3BOT*0.5*(T3BOT-Tamb));

qTOT = qTOP+qBOT+qLAT;

```

Correlated Weighted Heater Clamp Convective Loss Calculator

```

function [qTOP,qBOT,qLAT,qTOT] = qHCweighted(Tsurf,Tamb,Cweight)

L = 0.017;
W = 0.032;
H = 0.014;

%Top and Bottom Surface Geometry
A_TB = L*W;
P_TB = 2*(L+W);
Lchar_TB = A_TB/P_TB;

```

%Top Surface Losses

```
Tcent1 = Tsurf;  
Tedge1 = 0.7159*Tsurf+8.4628;  
Tfcent1 = Tfilm(Tcent1,Tamb);  
Tfedge1 = Tfilm(Tedge1,Tamb);  
Ra1C = Ra(Tcent1,Tamb,Lchar_TB);  
Ra1E = Ra(Tedge1,Tamb,Lchar_TB);  
Nu1C = NuTOP(Ra1C);  
Nu1E = NuTOP(Ra1E);  
h1C = hCONV(Nu1C,Kfluid(Tfcent1),Lchar_TB);  
h1E = hCONV(Nu1E,Kfluid(Tfedge1),Lchar_TB);  
qTOP = A_TB*(h1C*Cweight*(Tcent1-Tamb)+h1E*(1-Cweight)*(Tedge1-Tamb));
```

%Bottom Surface Losses

```
Tcent2 = 0.8939*Tsurf+5.4790;  
Tedge2 = 0.5722*Tsurf+16.4191;  
Tfcent2 = Tfilm(Tcent2,Tamb);  
Tfedge2 = Tfilm(Tedge2,Tamb);  
Ra2C = Ra(Tcent2,Tamb,Lchar_TB);  
Ra2E = Ra(Tedge2,Tamb,Lchar_TB);  
Nu2C = NuBOT(Ra2C);  
Nu2E = NuBOT(Ra2E);  
h2C = hCONV(Nu2C,Kfluid(Tfcent2),Lchar_TB);  
h2E = hCONV(Nu2E,Kfluid(Tfedge2),Lchar_TB);  
qBOT = A_TB*(h2C*Cweight*(Tcent2-Tamb)+h2E*(1-Cweight)*(Tedge2-Tamb));
```

%Lateral Surfaces Losses

```
A_LAT = 2*(L*H+W*H);  
TlatTOP = 0.7967*Tsurf+7.9244;  
TlatBOT = 0.7172*Tsurf+9.5940;  
TflatTOP = Tfilm(TlatTOP,Tamb);  
TflatBOT = Tfilm(TlatBOT,Tamb);  
Lchar3 = H;  
A3 = 2*(L*H+W*H);  
Ra3TOP = Ra(TlatTOP,Tamb,Lchar3);  
Ra3BOT = Ra(TlatBOT,Tamb,Lchar3);  
Nu3TOP = NuLAT(Ra3TOP,Pr(TflatTOP));  
Nu3BOT = NuLAT(Ra3BOT,Pr(TflatBOT));  
h3TOP = hCONV(Nu3TOP,Kfluid(TflatTOP),Lchar3);  
h3BOT = hCONV(Nu3BOT,Kfluid(TflatBOT),Lchar3);  
qLAT = A_LAT*(h3TOP*0.5*(TlatTOP-Tamb)+h3BOT*0.5*(TlatBOT-Tamb));  
  
qTOT = qTOP+qBOT+qLAT;
```

Device Conductivity Calculator (Single Value Input)

```
function [kEFF] =  
conductivity(TBox1,TBox2,Tline1,Tline2,Tamb,V,I,deltaX,Cweight)  
  
A = 1.68E-5;  
  
P = V*I;  
deltaT = Tline1 - Tline2;  
dTdx = deltaT/deltaX;  
  
[qTOP_HC qBOT_HC qLAT_HC qTOT_HC] = qHCweighted(TBox1,Tamb,Cweight);  
[qTOP_DC qBOT_DC qLAT_DC qTOT_DC] = qDCweighted(TBox2,Tamb,Cweight);  
qLOST = qTOT_HC + qTOT_DC;  
  
kEFF = (P-qLOST)*deltaX/(A*deltaT);
```

Device Conductivity Calculator (Data Matrix Input)

```
function [kEFF] = conductivityVECT(Data,Tamb,Cweight)  
  
deltaX = 0.003;  
A = 1.68E-5;  
  
V = Data(:,1);  
I = 0.001*Data(:,2);  
TBox1 = Data(:,3);  
TBox2 = Data(:,4);  
TLine1 = Data(:,5);  
TLine2 = Data(:,6);  
  
for n = 1:length(Data)  
  
    P(n) = V(n)*I(n);  
    deltaT(n) = TLine1(n) - TLine2(n);  
    dTdx(n) = deltaT(n)/deltaX;  
  
    [qTOP_HC(n) qBOT_HC(n) qLAT_HC(n) qTOT_HC(n)] =  
        qHCweighted(TBox1(n),Tamb,Cweight);  
  
    [qTOP_DC(n) qBOT_DC(n) qLAT_DC(n) qTOT_DC(n)] =  
        qDCweighted(TBox2(n),Tamb,Cweight);  
  
    qLOST(n) = qTOT_HC(n) + qTOT_DC(n);  
  
    kEFF(n) = (P(n)-qLOST(n))*deltaX/(A*deltaT(n));  
  
end
```