Silicon Cast Wafer Recrystallization for Photovoltaic Applications

By

Eerik T. Hantsoo
B.S., Mechanical Engineering
Stanford University, 2005

Submitted to the Department of Mechanical Engineering in Partial Fulfillment of the Requirements for the Degree of

Master of Science
at the
Massachusetts Institute of Technology

June 2008

Signature of Author: _____
Department of Mechanical Engineering
May 21, 2008

Certified by: _____
Emanuel M. Sachs
Fred Fort Flowers and Daniel Fort Flowers Professor of Mechanical Engineering
Thesis Supervisor

Accepted by: _____
Lallit Anand
Professor of Mechanical Engineering
Chairman, Committee for Graduate Students

Silicon Cast Wafer Recrystallization for Photovoltaic Applications

By

Eerik T. Hantsoo

Submitted to the Department of Mechanical Engineering
on May 21, 2008 in Partial Fulfillment of the
Requirements for the Degree of Master of Science

ABSTRACT

Current industry-standard methods of manufacturing silicon wafers for photovoltaic (PV) cells define the electrical properties of the wafer in a first step, and then the geometry of the wafer in a subsequent step. The geometry is typically defined by a combination of grinding and abrasive wire sawing. While mature, these processes remain slow and wasteful of raw materials. As the PV industry scales to meet increasing global demand for renewable energy, new processes for creating wafers must be explored.

This project sets out to enable high-speed casting of individual wafers, by developing a zone recrystallization process to improve rapid-cast wafers of low electrical quality. In the process, individual wafer geometry is defined in an upstream high-speed casting step with little regard to electrical quality. Subsequently, the electrical properties (through grain structure, dislocation density, and segregation of impurities) are optimized by zone recrystallization.

The work outlined in this report documents the development of a custom, high-purity zone recrystallization furnace; an encapsulation mechanism for molten wafers; a mechanical fixturing scheme to preserve the planarity of recrystallized samples; and a release layer to prevent adhesion of the wafer to support structures.

Further, the results of experiments investigating temperature profile effects on defect density and grain structure are discussed. Specifically, results demonstrating completely redefined grain structure and improved dislocation density are disclosed. Minority carrier lifetime measurements are also disclosed. Although still preliminary, overall results are promising for the successful refinement of small-grained, rapid-cast wafers into large-grained, high-lifetime wafers suitable for use as high-efficiency PV cells.

Thesis Supervisor: Emanuel M. Sachs
Title: Fred Fort Flowers and Daniel Fort Flowers Professor of Mechanical Engineering

Acknowledgements

I would like to thank my advisor, Prof. Ely Sachs, for supporting me throughout my graduate studies. Without his patience, guidance, and expansive cross-disciplinary knowledge, this project would not have been possible.

I would like to thank several other members of the Sachs Group whose contributions to this project were indispensable. Jim Serdy was responsible for a large share of furnace development and construction, as well as much of the early experimental work detailed in this document. His design and fabrication efforts have been at the core of this project from the start. Laura Zaganjori's efforts have been pivotal in ensuring consistent progress (for me and many others in the lab) over the course of this project. Anjuli Appapillai, Chris Ruggiero, and Alison Greenlee, my labmates within the cast wafer project, have provided support in more ways than I can count. Likewise, my thanks go out to all other students and staff within the Sachs Group for their help and feedback.

Beyond the Sachs Group, I would like to thank some specific members of the Mechanical Engineering community: Aaron Gawlik for his support in programming microcontrollers for pyrometer control, Shawn Chester for his help in understanding the underlying solid mechanics, and Prof. Tonio Buonassisi for providing thorough explanations of solar cells at the device level.

I would like to thank Doug Spreng for his generous gift to the project, which helped to support my Research Assistantship. I would also like to thank the Fred Fort Flowers and Daniel Fort Flowers Chair for support of this work.

I would like to thank my undergraduate advisor, Prof. Beth Pruitt, for opening the door to academic research to me, as well as for introducing me to Ely. Finally, and most of all, I would like to thank my family and friends for (in each their own way) putting me on the path I'm on. I wouldn't have made it without you.

Table of Contents

Chapter 1: Background

Generating economical, scalable, renewable energy in globally relevant quantities is among the foremost challenges facing the science and technology communities in the coming decades. Of the sources of renewable energy currently available, solar has the largest total flux and resource potential.[1]

1.1 Energy Demand

Current Demand and Growth

As of 2005, worldwide energy demand is 488 exajoules (EJ) per year, which corresponds to a steady-state power draw of approximately 15 terawatts (TW).[2] Of the global energy mix, over 86% derives from fossil fuels and 6% derives from nuclear power.[3] The remainder derives from various renewable sources.

United States energy demand has grown fairly steadily at a 2% combined annual growth rate over the past 55 years.[4] This growth rate maps to a demand of 240EJ in the year 2050. Globally, energy demand has risen at a combined annual growth rate of 2.2% over the past 25 years, which maps to a global demand of 1180EJ in the year 2050.[5] If the current distribution of energy sources persists to that point of energy use, resource scarcity and strains on the environment may come to bear in detrimental and far-reaching ways.

Solar Resource Base and Renewable Competitors

Of the sources of energy incident on earth, the total flux of solar energy by far the largest. 89,000 TW of solar energy are incident, with 30% of that power incident on dry land. Assuming only land-based collection, 1800 times the current global energy demand

is incident in the form of solar power. Other renewable sources offer promise (for instance, wind at 100 TW global resource capacity, waves at 200 TW, photosynthesis for biomass at 30 TW), but not the vast scale that the solar resource offers.

Table 9.2. Worldwide Resource Bases for Renewable Energy: Fluxes and Stored Energy*

Resource	Energy Type	Total Flux TW (quads/yr)	Net Stored Energy TW-yr (quads)**
Solar	land flux	27,000	0
	ocean flux	53,000	0
	total flux	89,000 (2.7×10^6)	0
Wind	kinetic energy (land)	100 (3,000)	0
Waves	kinetic energy (ocean), gravity	200 (6,000)	0
Biomass	photosynthesis	30 (900)	750 (22,400)
Hydro	latent evaporative heat	40,000 (1.2×10^6)	80,000 (2.4×10^6)
Tides	gravitational forces between the moon and earth	> 3 (90)	3 (90)
Geothermal			
– hydrothermal – geopressured and – magma	convection	~2.5 (75)	1.710^5 (5.1×10^6)
– hot dry rock	conduction	≥ 30 (900)	$3.3\ 10^6$ ($1\times\ 10^8$)
All fossil fuels	chemical reaction enthalpies or heats ***	10 (300)	≥ 360,000 (1.1×10^7)

*estimates based on information provided by Armstead and Tester (1987), Grubb (1990), and Armstead (1983); US figures based in part on USGS estimates (e.g., see Muffler and Guffanti (1978), Sass (1993)).
**1 quad = 3.35×10^{-2} TW-yr = 10^{15} Btu
***does not include any energy stored from the solar greenhouse effect, e.g., in atmospheric CO_2, H_2O, or CH_4.

Figure 1: From *Sustainable Energy: Choosing Among Options*, by Tester, Golay, Drake, Driscoll, and Peters. Solar is by far the largest incident energy flux.[6]

Nuclear power represents a known means of generating large amounts of baseload energy with relatively little lifecycle greenhouse gas emission; however, a reliable method of waste disposal has yet to be established, particularly given the long half-lives of hazardous radionuclides present in nuclear waste. Furthermore, proliferation concerns,

lengthy licensure processes, and supply chain bottlenecks continue to stymie the large-scale implementation of nuclear energy in many localities.

Although the extant resource base of fossil fuels is large (360,000 TW-yr), the effects of continuing their use as an energy source are unacceptable, given the critical consequences of CO_2-induced climate change. Atmospheric CO_2 levels have risen from 280ppm in 1750 to 377ppm in 2008.[7] Furthermore, the CO_2 concentration growth rate has been largest in the years from 1995-2005 (1.9ppm per year) as opposed to the average since measurements have been taken 1960-2005 (1.4ppm per year).[8] This acceleration in CO_2 concentration has been caused primarily by increased energy generation.[9]

The consequences of continued fossil fuel burn at the current rate are significant. Anthropogenic climate change has the potential to disrupt food supply, alter weather patterns, raise sea levels, damage ecosystems, and decrease fresh water supplies, among many others.[10] Thus, less carbon-intensive sources of energy remain the most promising option for mitigating the negative effects of climate change. As such, solar energy offers the greatest potential for minimal lifecycle carbon emissions and globally relevant scale.

1.2 Solar Market Penetration

Current Extent

In 2006, solar energy sources supplied less than 0.1% of both US and worldwide energy.[11] Of worldwide photovoltaic (PV) device production in 2006, 90% was wafer-based silicon (as opposed to thin-film).[12] Thus, although very little installed solar capacity exists as a share of overall energy production, the large majority of devices now being manufactured are based on silicon wafers. Consequently, innovations in silicon

wafer-based solar cells represent a significant opportunity to impact a large segment of an industry with huge growth potential.

Reasons for Lack of Penetration

Historically, the high levelized costs of PV energy have limited its penetration into global energy markets. Levelized energy cost of installed PV is highly dependent on the location of the installation, balance of system costs, and the discount rate employed in analyzing future energy flows. Values typically range from $0.20/kWh in commercial installations to $0.40/kWh in residential installations.[13] Although government subsidies on PV in different localities have decreased the effective levelized cost, unsubsidized PV energy is more expensive than grid electricity in all but a few OECD nations and all but a few local regions of the US.[14]

Capital costs tend to dominate the up-front costs of installing PV, particularly as installation size scales up. In wafer-based modules, the cost of the wafer typically constitutes 40-45% of the total module cost, which is the largest single contributor to up-front capital costs.[15, 16] Thus, by decreasing the cost of the wafer, potentially large fractions of total installed cost can be decreased, in turn driving down levelized PV energy costs and driving up market penetration.

1.3 Current Wafer Manufacturing Practices

Multicrystalline Silicon

Block casting followed by directional solidification is the primary means of crystal growth used for multicrystalline wafers. In this technique, a fused quartz crucible (with a Si_3N_4 powder release liner) is filled with solar-grade silicon feedstock (usually

small-diameter microcrystalline chunks) from the Siemens process, melted, and allowed to cool back down into a solid block. These blocks are large, typically on the order hundreds of kilograms. By imposing directionality (bottom to top) on the freeze front, moving the freeze front slowly through the bulk of the block, and minimizing potential sources of nucleation, large grain sizes and low dislocation densities are achievable.

Following solidification and cooling, the block is sawn down into smaller blocks, which in turn are wire sawn down into individual wafers. These sawing processes typically use abrasive wire cutting to generate wafers. In the drive toward greater silicon feedstock utilization, recent trends have favored increasingly thin wafers; wafers below 200μm are now common within the PV sector.

As wafers grow thinner, kerf losses from the wire sawing process constitute an increasing fraction of non-utilization of silicon feedstock. In some cases, the amount of silicon lost to kerf actually exceeds the amount of silicon employed in usable wafers.[17] Worse still, silicon from kerfs is difficult to reclaim, as it tends to have high transition metal content – the wires used for sawing are steel, and some metal particles invariably make it into the sawing waste stream.

Single Crystal Silicon

Single crystals of silicon are typically grown in a Czochralski crystal grower. After being seeded by a single crystal fragment of silicon, a cylindrical boule is slowly extracted from a melt of silicon in a fused quartz crucible while being rotated. As no solid surfaces are in contact with the growing boule (save the seed crystal), no new nucleation sites are present and growth is single-crystal. An important constraint of this type of single-crystal growth is that boules are necessarily cylindrical, which dictates that

(given a certain minimum device packing density in a completed module) some of the perimeter of the boule must be ground off, in order to yield a rectangular wafer.

In the same way that multicrystalline wafers are wire sawn from cast blocks, single crystal wafers are wire sawn from boules, presenting the same silicon utilization issues.

Drawbacks of Conventional Wafer Manufacturing Methods

Current methods of manufacturing wafers are based heavily on hardware and techniques borrowed from the microelectronics industry, where overall trends have tended toward compaction. As dies of increasing complexity and functionality are packed into the same small area on a wafer, development and fabrication costs increasingly dominate overall cost structures of microelectronics manufacturing (as opposed to the costs of raw materials).[18]

While PV shares the same base material with microelectronics, the scaling challenge is fundamentally different. Rather than decreasing device size and packing more devices onto a wafer, the challenge in photovoltaics is to produce relatively simple, large-area devices as cheaply as possible. There is no benefit to in-plane miniaturization. As such, costs and throughput of the feedstock – wafers – are of much greater cost significance in PV than in microelectronics. Techniques such as wire sawing are economically viable in microelectronics, as product prices corresponding to $\$100/cm^2$ at the wafer level are profitable.[19] In PV, product prices must fall well below $\$0.01/cm^2$ at the wafer level in order to be cost-competitive with other sources of energy. Accordingly, low-throughput and resource-inefficient process steps should be eliminated if possible.

Wire sawing consumes large amounts of wire (which is used only once) and slurry, wastes a significant fraction of feedstock material, and typically has a long Takt time. As such, the process is ripe for improvement or replacement.

1.4 This Work: A New Approach

Process Outline

Wire sawing is demanded by the constraint that defining the geometry of the wafer takes place after the grain structure has been defined. Not only must solidification be carried out slowly, but subsequent definition of wafer geometry must be carried out on a solid, brittle material – necessitating the slow, abrasive process of wire sawing.

By defining the grain structure (and the resulting electrical properties) *after* the geometry of the wafer has been defined, wire sawing can be eliminated altogether. This work seeks to enable high-speed casting methods for creating individual wafers (thereby eliminating wire sawing) by investigating the subsequent refinement of material structure for electrical performance.

Above a certain maximum casting speed, high-speed casting methods necessarily yield small-grained material with low minority carrier lifetime.[20] Although the economics of high-speed wafer casting (conceivably, similar to those of injection molding, die casting, or centrifugal casting) are compelling, resulting material quality has been sufficiently poor to date that the technique has not been adopted. Industry has been reluctant to take on the challenge of refining the electronic quality of high-speed cast wafers: wafer-by-wafer refinement is required, as opposed to refinement of an entire mass at once. As such, issues such as wetting, contamination, and geometric constraint

must be dealt with on individual wafers, potentially cutting into yields. Surface effects are much more prevalent per unit product in a wafer than in an ingot. In this way, a key enabling innovation that must be implemented is an effective means of encapsulating the wafer, to prevent contamination, wetting, and balling.

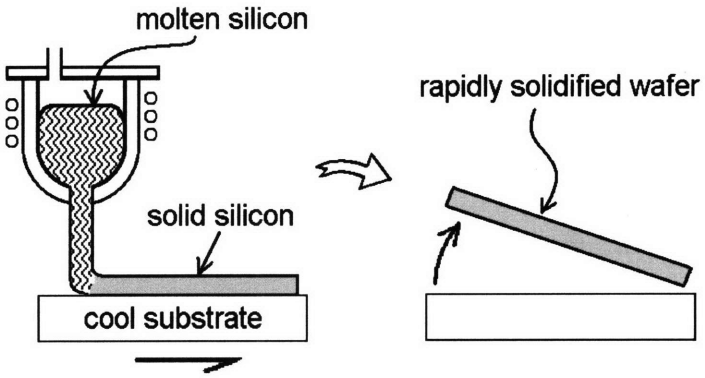

Figure 2: Schematic of high-speed casting process desired as an alternative to block casting and sawing.

This works seeks to develop a suitable post-casting material refinement step. The refinement consists of growing an oxide capsule around a cast wafer, then zone melting the wafer and recrystallizing it in a controlled manner. With a successful refinement, the dislocation density, purity, and minority carrier lifetime of high-speed cast wafers would equal that of conventional sawn wafers, yielding performance-competitive solar cells at a lower cost.

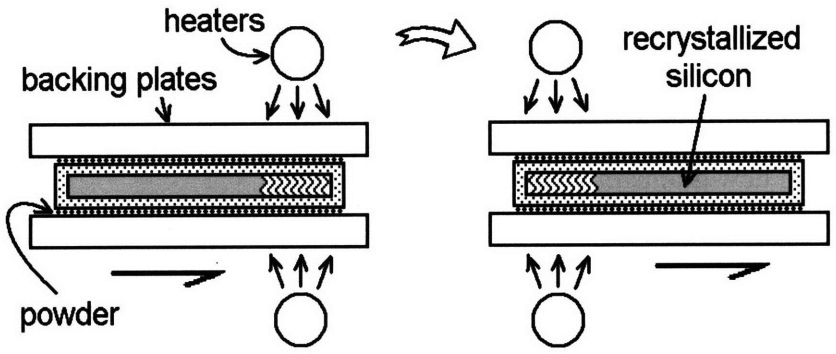

Figure 3: Conceptual schematic of the recrystallization process investigated in this work. The powder functions as a release layer, to prevent sticking of the oxide capsule to the backing plates (see 3.2).

While a multitude of parameters in this process will influence the resulting material quality, this work focuses mainly on the effects of temperature profile on the dislocation density and minority carrier lifetime of recrystallized silicon. These factors have a strong impact on completed device performance.

1.5 Summary of Progress and Implications

Progress

In accordance with the motivations outlined thus far, this project succeeded in developing both the process and hardware necessary to recrystallize wafers at lab scale. The project furthermore demonstrated promising results, indicating that prospects are good for improving the electrical properties of small-grained wafers. Several parallel components had to function in order to enable these developments:

- A high-purity custom furnace was designed, built, and successfully operated reliably enough to recrystallize over 30 samples. This furnace is the core piece of

hardware for the process; it is what actually melts the silicon for recrystallization and controls the freeze front.

- A functional capsule to contain the wafer was demonstrated, with sufficient robustness to hold together at temperature and sufficient smoothness to avoid nucleating new grains in the pattern of old grains.
- A release layer was developed which reliably prevents adhesion of recrystallizing wafers to backing and capping plates.
- A scheme of geometrical constraint for the recrystallized wafer was devised which – although in need of improvement for reliability – succeeded in delivering a sample with less than 10% overall thickness variation.

Noteworthy results from experiments included:

- Complete melting and regrowth of totally new grain structures was demonstrated in recrystallized wafers, with grain boundaries principally out of plane of the wafer.
- A decrease in dislocation density of recrystallized samples was demonstrated relative to control (non-recrystallized) samples.
- Minority carrier lifetimes in recrystallized samples, which – while degraded from the level of controls – still resided within the same order of magnitude as control samples. This bodes well for future improvements.

Implications

While the high-speed, low-cost casting of silicon wafers will require significant further work, the results demonstrated in this document show that *refining* high-speed cast wafers has some likelihood of succeeding (at least at the technical, if not economic,

level). As refinement of electronic properties is among the most significant gating functions for exploring high-speed cast wafers, the results from this project lay early groundwork for eventual further expansion into cast wafers.

Cast wafers, in turn, present the possibility of significantly altering the cost structure and manufacturing process for wafer-based PV devices across the industry. Given growth in the PV industry and increasing global energy demand, making low-cost wafers is a potentially high-impact and lucrative contribution to the crystalline silicon photovoltaics industry at large.

References

---

[1] Tester, Golay, Drake, Driscoll, and Peters. *Sustainable Energy: Choosing Among Options*. Cabmbridge, MA: The MIT Press, 2005. 408-416.

[2] Energy Information Administration, U.S. Department of Energy. "International Total Primary Energy Consumption and Energy Intensity: Total Primary Energy Consumption (Quadrillion Btu)." 21 May 2008. <http://www.eia.doe.gov/emeu/international/energyconsumption.html>.

[3] Energy Information Administration, U.S. Department of Energy. "International Total Primary Energy Consumption and Energy Intensity: Primary Energy Consumption by Source [Energy or Fuel Type], 1980-2005 (Quadrillion Btu, Except as Noted)." 21 May 2008. <http://www.eia.doe.gov/emeu/international/energyconsumption.html>.

[4] Energy Information Administration, U.S. Department of Energy. "Table 1.1: Energy Overview, Selected Years 1949-2006." 21 May 2008. <http://www.eia.doe.gov/emeu/aer/pdf/pages/sec1_5.pdf>.

[5] Energy Information Administration, U.S. Department of Energy. "Table 11.1: World Primary Energy Production by Source, 1970-2004." 21 May 2008. <http://www.eia.doe.gov/aer/pdf/pages/sec11_3.pdf>.

[6] Tester, Golay, Drake, Driscoll, and Peters. *Sustainable Energy: Choosing Among Options*. Cambridge, MA: The MIT Press, 2005. 409.

[7] Blasing, T.J. and Karmen Smith. Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory. 21 May 2008. <http://cdiac.ornl.gov/pns/current_ghg.html>.

[8] Intergovernmental Panel on Climate Change Report, 2007. *Climate Change 2007: Synthesis Report*. Valencia, Spain, 2007. 37.

[9] Ibid.

[10] Intergovernmental Panel on Climate Change Report, 2007. *Climate Change 2007: Synthesis Report*. Valencia, Spain, 2007. 48.

[11] Energy Information Administration, U.S. Department of Energy. "U. S. Energy Consumption by Energy Source." 21 May 2008. <http://www.eia.doe.gov/cneaf/solar.renewables/page/trends/table1_1.html>.

[12] "The Q Factor, Sharp and the Market; Cell Technology Shares (%)." *Photon International*, March 2008. 152.

[13] Solarbuzz, LLC. "Solar Electricity Prices, May 2008." 21 May 2008. <http://www.solarbuzz.com/SolarPrices.htm>.

[14] Energy Information Administration, U.S. Department of Energy. "State Electricity Profiles." 21 May 2008. <http://www.eia.doe.gov/cneaf/electricity/st_profiles/e_profiles_sum.html>.

[15] Swanson, Richard. Sunpower Corporation. "Cost Reduction Potential of Wafered Silicon PV." 21 May 2008. <http://www.ases.org/solar2007/presentations/wednesday/1030am/forums/1-pv/2-swanson.pdf>.

[16] Rogol, Michael. "Refining Benchmarks and Forecasts; Integrated Cost of C-Si Module – sum-of-best practices." *Photon International*, February 2008. 92.

[17] Ibid.

[18] Arvind, Massachusetts Institute of Technology. "Do We Need More Chips (ASICs)?" 21 May 2008. <http://csg.csail.mit.edu/6.375/handouts/lectures/L01-Introduction.pdf>.

[19] Intel Corporation. "Intel Details Upcoming New Processor Generations." 21 May 2008. <http://www.intel.com/pressroom/archive/releases/20070328fact.htm>.

[20] Belouet, C. "Growth of Silicon Ribbons for Terrestrial Solar Cells by the RAD Process." *Silicon Processing for Photovoltaics I.* Ed. Khattak and Ravi. Amsterdam: Elsevier, 1985. 106-110.

Chapter 2: Theoretical Basis

Imperfections in silicon can be grouped into three categories: point (0-dimensional), line (1-dimensional), and sheet (2-dimensional) defects.[1] While each are addressed within this work, particular focus has been placed on both line and sheet defects.

2.1 Point Defects

Point defects refer to vacancies, interstitials, or impurities in a crystal lattice. Although these sources of crystal imperfection were considered in this work (as in, for instance, the avoidance of transition metal impurities in all hardware components), the primary focus was not on point defects.

The reason for focusing on line and sheet defects is that the concentration of point defects cannot be reduced below some baseline level in a real crystal. Once point defects have been frozen into a crystal (after the initial liquid to solid phase transformation), a certain minimum concentration of point defects is in thermodynamic equilibrium within the crystal.[2] Thus, subsequent treatments are capable of only limited effect on the point defect concentration.

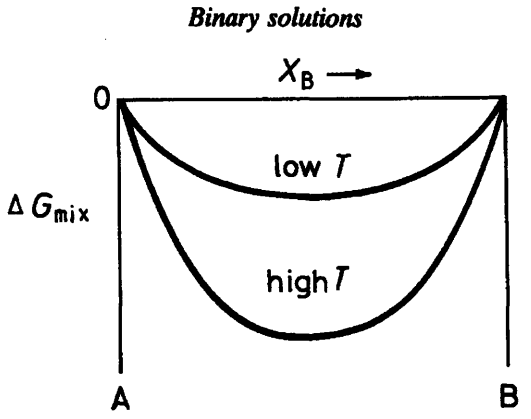

Fig. 1.9 Free energy of mixing for an ideal solution.

Figure 1: From Porter and Easterling's *Phase Transformations in Metals and Alloys*, fig. 1.9. In order to achieve complete purity of either A or B, the change in free energy becomes asymptotically large as purity increases.[3]

From a free energy standpoint: the Gibbs free energy change ΔG of adding a point defect to a crystal is negative.[4] As such, a crystal with some point defects is actually more thermodynamically stable than one without (due to the magnitude of entropy increase associated with the addition of a point defect). Thus, any real crystal above absolute zero will necessarily contain some point defects.

A final mechanism worth noting is that of zone purification. Zone purification occurs as a side effect of the recrystallization process investigated in this work, and it decreases the density of point defects: specifically, of impurities. Although unrelated to temperature profile, zone purification occurs naturally in settings with a moving solid-melt interface. In the case of this work, the driving force for zone purification is a difference between the liquid and solid solubility of a particular species in silicon. This difference is quantified by the segregation coefficient k:

$$C_S = kC_L$$

where C_S is the equilibrium impurity concentration in the solid, C_L is the equilibrium impurity concentration in the liquid, and k is the segregation coefficient.[5]

Table 4
Segregation coefficients of impurities in silicon.

Element	Segregation coefficient	Element	Segregation coefficient
Ag	1.7×10^{-5}	Mo	4.5×10^{-8}
Al	3×10^{-2}	Nb	$<4.4 \times 10^{-7}$
Au	2.5×10^{-5}	Ni	1.3×10^{-4}
B	8×10^{-1}	P	3.5×10^{-1}
C	5×10^{-2}	Pb	?
Ca	?	Pd	5×10^{-5}
Co	2×10^{-5}	Sn	3.2×10^{-2}
Cr	1.1×10^{-5}	Ta	2.1×10^{-8}
Cu	8.0×10^{-4}	Ti	2.0×10^{-6}
Fe	6.4×10^{-6}	V	4×10^{-6}
Gd	$<4.0 \times 10^{-7}$	W	1.7×10^{-8}
Mg	3.2×10^{-6}	Zn	1.0×10^{-5}
Mn	1.3×10^{-5}	Zr	$<1.6 \times 10^{-8}$

Figure 2: Segregation coefficients of commom silicon impurities, from Wang, Khattak, and Ravi: *Silicon Processing for Photovoltaics I.* Harmful impurities such as zirconium and iron have strong zone refining potential; boron and phosphorous do not.[6]

Presuming a given impurity species has a higher solubility in liquid silicon than in solid silicon, the concentration of that impurity will increase in the melt; it will take up rejected impurity from the solidifying front of silicon.

When this freeze front is moving, the overall effect is that impurity concentration increases in the moving melted portion, and is left lower throughout the rest of the sample. At the conclusion of recrystallization, the residual bump of silicon which ejects out the trailing edge of the sample (the last part of the sample to freeze) – in theory – should have high levels of impurities relative to the rest of the wafer. Although this phenomenon has not been tested specifically as it relates to this process, future work could establish the extent of beneficial effects.

As wafers entering the zone recrystallization process will be doped with either phosphorous or boron, it is important to note that the segregation coefficients of these materials are weak. Thus, doping levels will be minimally affected by zone recrystallization.

2.2 Line Defects

Line defects refer to dislocations, which are linear discontinuities in a crystal structure.[7] There are two types of dislocations: edge and screw. Edge dislocations are incomplete crystal planes, which induce stress in neighboring planes of atoms. Screw dislocations are a discontinuity and effective shearing of several stacked planes of atoms.

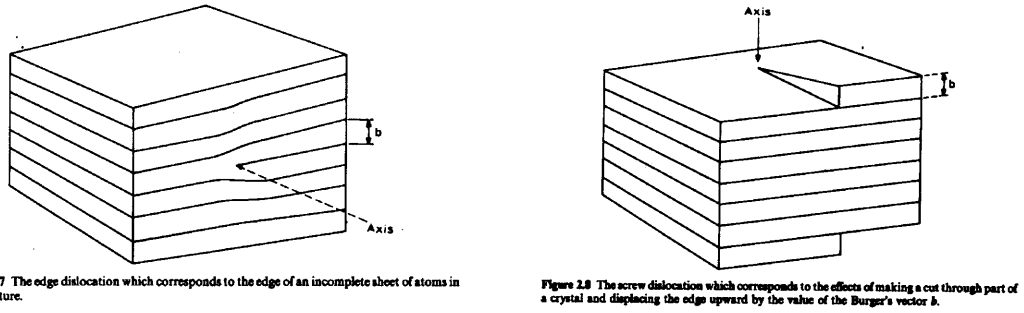

Figure 2.7 The edge dislocation which corresponds to the edge of an incomplete sheet of atoms in the structure.

Figure 2.8 The screw dislocation which corresponds to the effects of making a cut through part of a crystal and displacing the edge upward by the value of the Burger's vector *b*.

Figure 3: Edge (left) and screw (right) dislocations, as depicted in *Crystal Growth Processes* by J. C. Brice. [8]

Unlike point defects, dislocations are not in thermodynamic equilibrium in the bulk of a crystal; they can be annealed out over time due to a net repulsive force between dislocations of the same sign.[9] Dislocations in crystalline silicon solar cells are detrimental to device performance, as dislocations act as recombination centers for mobile charge carriers. Solar cells made from highly-dislocated silicon suffer from reduced minority carrier lifetime, and consequently lower overall efficiency. Hence, one of the goals of the present work is to minimize dislocation density.

Dislocation Formation

Dislocations form primarily as a result of thermoelastic shear strains present during crystal growth and subsequent cooling.[10] Given thermally-induced strains, dislocations will form in order to relieve these strains. Dislocations can also be initiated by mid-grain point defects acting as Frank-Reed sources (which speaks to the importance of minimizing contaminants); however, this work will focus only on dislocations caused by thermoelastic shear strains.[11]

Second Derivative Dependence of Stress on Temperature Profile

In a cooling crystal, the formulation for stress-free conditions is

$$\nabla^2 T = \frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2} + \frac{\partial^2 T}{\partial z^2} = 0 . \quad [12]$$

Samples subject to a temperature profile in any direction experience thermal stress proportional to the second partial derivative of temperature in any given direction. In a flat plate, such as the wafers used in this work, the thermal stress σ_{th} scales as

$$\sigma_{th} \propto \frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2} . \quad [13]$$

The occurrence of this stress can be qualitatively explained by envisioning a wafer under two different temperature profiles: one linear and the other curved.

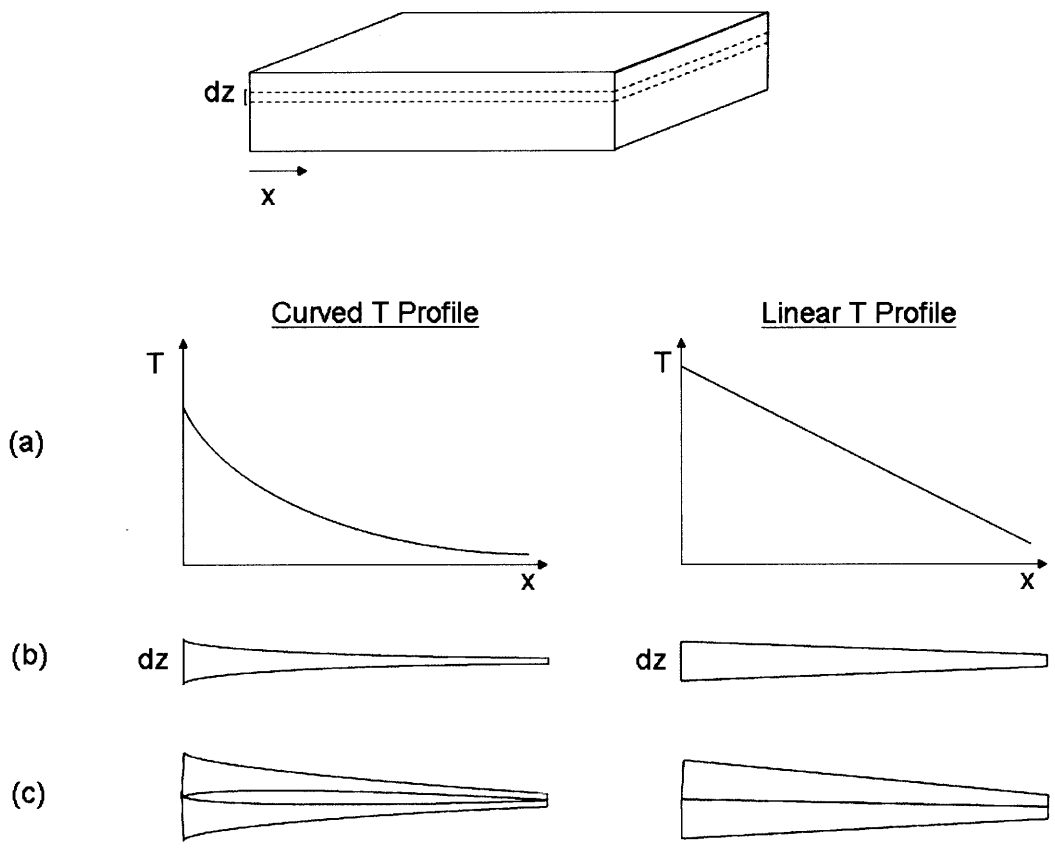

Figure 4: Qualitative representation of the effects of curved and linear temperature profiles on stresses in a wafer. The wafer is subjected to a curved temperature profile (on the left) and a linear temperature profile (on the right), as shown in (a). Given a differential slice of the wafer thickness dz, thermal expansion will cause the thickness of the wafer to vary linearly (shown on the right) and nonlinearly (shown on the left), in row (b). When multiple slices dz of the wafer are stacked back together, as shown in row (c), the slices subjected to a nonlinear profile will not fit back together into a continuous solid unless they are stressed.

Given a certain ΔT across the length of a wafer, a linear temperature profile is required in order to minimize stress in the wafer.

The wafers recrystallized in this work are subject to steep temperature gradients upon cooling. Due to the need to reject heat conductively from the silicon freeze front (which requires a certain minimum $\frac{dT}{dx}$) as well as the boundary condition of equilibrating out to room temperature at the opposite end of the wafer, a perfectly linear temperature profile is impossible to achieve. Furthermore, recrystallizing at

progressively higher speeds requires that the temperature gradient at the freeze front is steeper still; the temperature profile deviates increasingly from linearity as pull speed increases. Consequently, all wafers that have been recrystallized are subject to thermoelastic shear stresses upon cooling. These stresses increase with increasing pull speed due to the increasingly steep temperature gradient required at the freeze front to conduct heat away.

Given a perfectly elastic, unbuckled wafer at all temperatures, thermal stresses would simply dissipate once the wafer had cooled uniformly to room temperature. However, silicon exhibits highly plastic behavior at elevated temperatures. In the plastic regime, stresses are readily relieved by the formation and migration of dislocations.

Plastic Flow

Silicon undergoes a brittle-ductile transition (BDT) at a temperature between 500°C and 1000°C.[14] Variability in measured values of BDT has to do with the measurement technique used, stress induced, and purity level of silicon used. BDT temperature corresponds to the point at which the energy of thermal excitation begins to exceed the energy required for the glide propagation of dislocations (glide is activated at lower temperatures, and so it – instead of climb – tends to dictate BDT).[15]

Dislocation motion is thermally activated and follows an Arrhenius relationship; dislocation mobility v scales as:

$$v \propto \tau^m \exp(\frac{-U}{kT})$$

where τ is the stress, U is the activation energy of dislocation motion, k is the Boltzmann constant, T is the absolute temperature, and m is an experimentally determined value.[16]

Above the BDT temperature, dislocations are free to migrate by this mechanism, relieving stress. At even higher temperatures – those approaching the melting point of silicon – dislocations are sufficiently mobile that stress relief is extremely fast, to the point that most dislocations cancel out before they are frozen in.[17]

2.3 Sheet Defects

Sheet defects are discontinuities that extend across many atomic planes; at the macroscopic scale, sheet defects are grain boundaries. Like dislocations, grain boundaries function as recombination centers in crystalline silicon solar cells; thus, it is advantageous to minimize grain boundaries (by increasing grain sizes). The role of grain boundaries as dislocation centers is manifest in higher average efficiencies of commercial single-crystal solar cells over multicrystalline cells.

Like dislocations, it is in practice possible to grow grain-boundary free macroscopic crystals; indeed, the entire microelectronics industry (and a subset of the solar cell industry) uses single-crystal silicon as a starting material. However, growth of grain-boundary free crystals is slow and consequently costly. The current preponderance of multicrystalline solar cells on the market represents a compromise between cost (via process simplicity and throughput advantages for multicrystalline silicon) and performance (via increased recombination in multicrystalline devices).

Grain Nucleation

New grains nucleate when a solid (whether a precipitate out of liquid or a feature on the capsule wall) in the solidifying melt exceeds the critical nucleation radius r^*.[18] Critical nucleation radius is determined by the tradeoff between solid-melt surface free

energy and the volumetric free energy change of solidification.[19] Dislocations and impurities in either the growing crystal or the physical envelope holding the melt can nucleate new grains, giving rise to multicrystalline growth. Presuming no features of critical nucleation radius are present in the volume of liquid immediately neighboring the freeze front, only existing grains will continue to grow.

Growth methods such as Czochralski (CZ) and float zone (FZ) benefit from the fact that the only solid in contact with the freeze front is the preexisting crystal. Presuming FZ or CZ growth is seeded with single crystal material, large single crystal boules can be grown. The process being developed in this work has the disadvantage that the oxide capsule required to prevent wetting of backing plates is in direct contact with growing grains at the freeze front; the capsule is capable of nucleating new grains. However, by optimizing certain aspects of grain growth, the negative impact of grain nucleation can be minimized.

Grain Boundary Orientation

From the standpoint of completed device performance, grain boundaries parallel to the plane of the wafer are most detrimental. Photogenerated charge carriers must find their way from the bulk of the wafer up to the junction (at the top of the cell); grain boundaries encountered at midplane cause these carriers to recombine before reaching the junction. Thus, while in-plane grain boundaries are not *intrinsically* more detrimental than out-of-plane grain boundaries, geometry renders them more damaging. Much more of the planview area of the cell is affected by an in-plane boundary than an out-of-plane boundary.

The orientation of grain boundaries can be controlled by temperature profile, pull rate, or by biasing the heat flow in the freezing wafer toward either the top or bottom of the wafer. Tilting away from vertical minimizes the impact of competing grains nucleating from both the top and bottom of the freeze front and colliding at midplane.

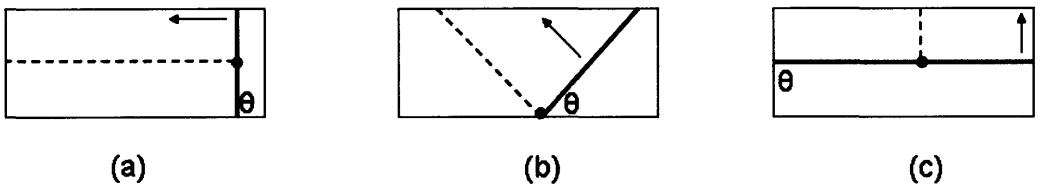

Figure 5: The effects of freeze interface tilt on grain boundary orientation. The freeze front is shown in bold in each sketch, propagating in the direction of the arrow. Given a grain boundary originating at the bold spot on each freeze front, the grain boundary (shown as a dashed line) will propagate roughly perpendicular to the freeze front. In (a), with a freeze interface angle $\theta = 90°$ to the plane of the wafer, the grain boundary grows parallel to the plane of the wafer. In (b) with $\theta = 45°$, the grain boundary terminates at the top surface of the wafer, causing less damage to device performance. Device performance is minimally damaged in case (c) where $\theta = 0°$.

Indeed, from the standpoint of minimizing mid-plane grain boundaries, a purely horizontal freeze front is desirable. As the present process does not allow horizontal growth (due to its continuous end-to-end nature), the ideal case is then to minimize θ.

In the ideal case, thermal symmetry and θ are interrelated. Their relationship can be explained as follows.

Heat of fusion rejection from the freeze front can be modeled as

$$\dot{Q}_{fusion} = wtv\rho H_{fusion}$$

where w is the width of the wafer, t is the thickness of the wafer, v is the pull speed, ρ is the density of silicon, and H_{fusion} is the enthalpy of fusion of silicon, on a per-mass basis. Given the dimensions of wafers and pull speeds employed in experiments on dislocation density, heat of fusion must be extracted at a rate of 3.6W in order to maintain equilibrium.

In order to sustain this rate of heat extraction (if backing plates are neglected), sufficient temperature drop and sufficient interfacial area must exist at the freeze front to conduct the heat of fusion away from the melt and into the solid. In reality, heat flows into and out of the backing plates are dominant; they are on the order of 1 kilowatt as compared to 3.6W of heat of fusion flow. Coupled with heater spacing, the thermal diffusivity of the backing plates overwhelmingly dominates the overall heat transfer taking place in the recrystallization furnace.

Thermal Asymmetry Requirement

In order to tilt the freeze front, overall heat flow must be biased in order to favor either the top or bottom plane of the wafer. This bias can be imposed on heat flows through the recrystallizing wafer in several ways; one way is to vary the vertical spacing between heater elements and the samples being processed.

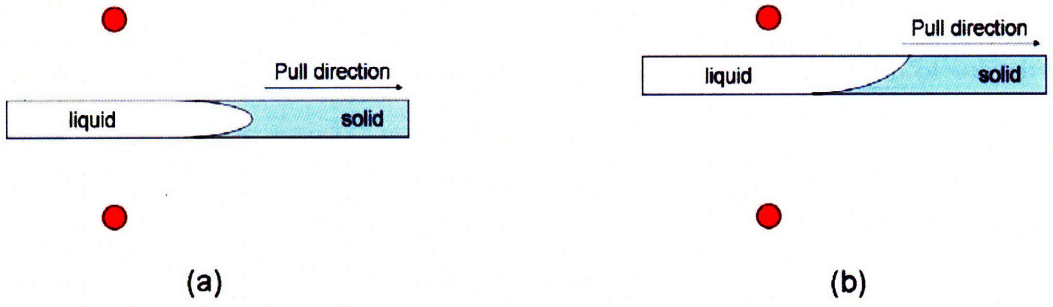

Figure 6: Neglecting backing plates, symmetric (a) and asymmetric (b) spacing of sample between heater elements (represented by the red circles). In (b), more heat is radiated into the top face of the wafer due to its proximity to the heater element; this bias causes the freeze front to migrate down the sample in the pull direction.

Another way to tilt the freeze front is to use backing plates of different thicknesses on the top and bottom of the wafer. Depending on the primary mode of heat

transfer through the backing plates – either radiation or conduction – the freeze front will tilt either up or down.

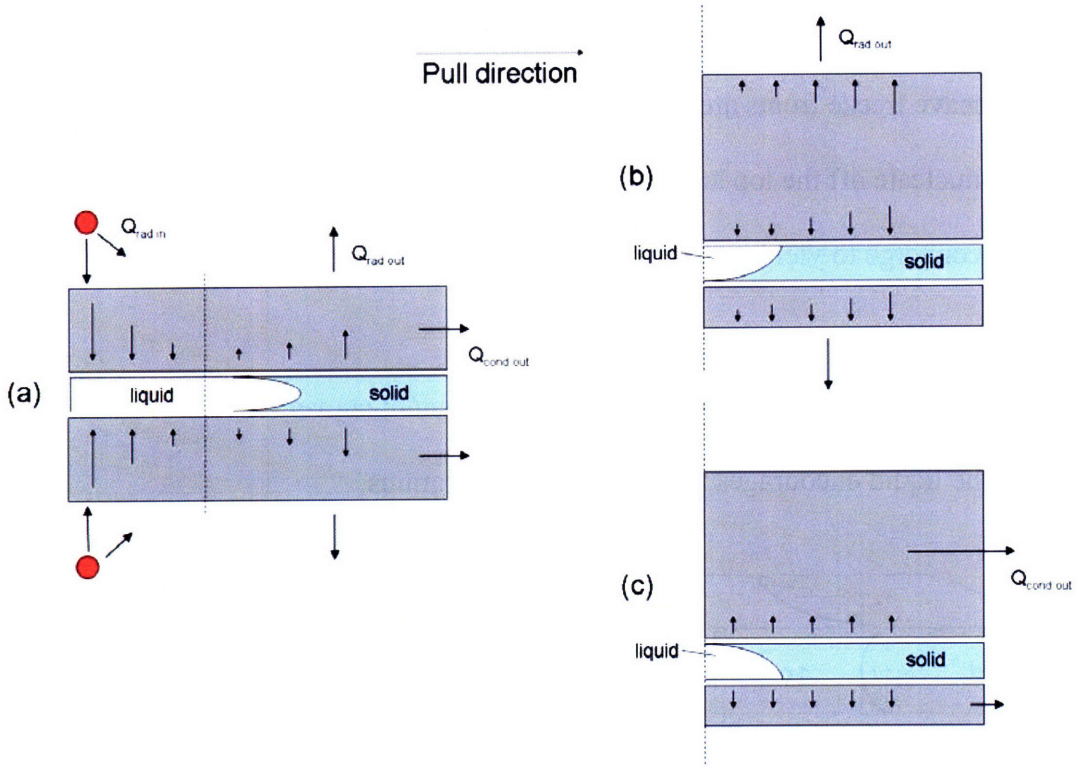

Figure 7: Three scenarios representing different heat transfer modes during freezing. (a) assumes that the backing plates both radiate and conduct heat effectively: heat flows away from the backing plate surfaces (once far enough removed from the heater elements; the inflection point is represented by the dotted line). This scenario results in a symmetric freeze front. (b) assumes the limiting case of dominant radiant cooling; the thermal mass of the upper backing plate must dissipate heat through the wafer, making the top side of the wafer hotter. (c) assumes the limiting case of dominant conductive cooling (through the backing plates); the larger cross-sectional area available to the cooling sample in the top plate makes the bottom of the sample hotter.

In the case of radiation-dominated cooling, the larger thermal mass of the top backing plate shifts the plane of thermal symmetry above the wafer entirely; this ensures that heat will principally flow down through the wafer. In the case of conduction-dominated cooling, the larger flux area available to the cooling sample through the thicker top plate makes the top of the wafer cooler, tilting the freeze front the opposite way. (As vertical adjustability of sample spacing was straightforward in the hardware of

this work, the simpler scheme of biasing sample spacing toward a heater was used when biased recrystallizations were desired.)

Freeze Front Curvature

Concave freeze fronts are suboptimal because they can lead to grain collisions; grains that nucleate off the top and bottom leading corners of the freeze front will eventually converge to yield a grain boundary parallel to the plane of the wafer.

Grain Divergence

A freeze front concave to the liquid is detrimental; conversely, a freeze front convex to the liquid encourages the growth of larger grains.

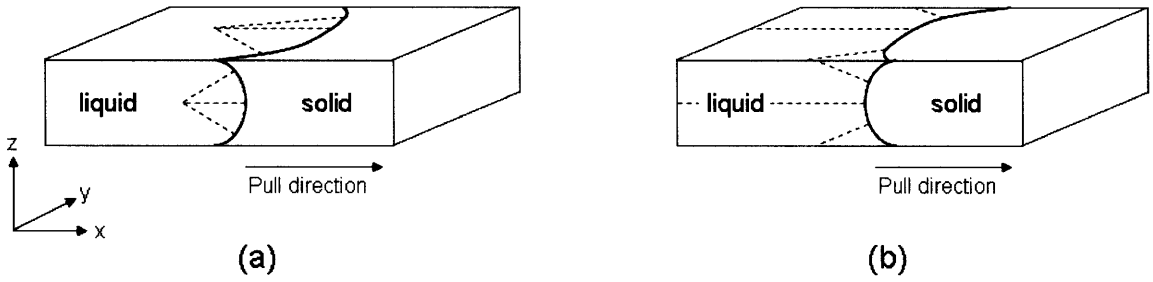

Figure 8: Non-ideal (a) and ideal (b) curvature modes of the freeze front. A freeze front concave to the liquid (as shown in (a)) promotes grain collision, whereas a front convex to the liquid (as shown in (b)) promotes grain divergence. Although preferred, the freeze front shape of (b) is impossible to achieve in practice using the present method, and is included purely for illustration.

Freeze front curvature in the y direction could presumably be controlled by employing backing plates of varying thickness: thinner or thicker in the middle of the y dimension than at the edges, depending on whether conduction or radiation dominates. This may be a useful area of future development.

Nucleation of Grains off Dislocations

Lattice distortions resulting from dislocations can reduce the free energy barrier required for the nucleation of new grains. Specifically, the interfacial free energy barrier

for nucleation is decreased in a dislocated lattice growth site; dislocated lattice sites are strained compressively.[20] By this mechanism, the required surface area for a critical nucleus is smaller and the overall free energy barrier for nucleation is lower. It is accordingly desirable to minimize dislocation density in the growing crystal, to prevent strained lattice sites at dislocations from causing new grains to nucleate.

The above factors (point, line, and sheet defects) all influence minority carrier lifetime, which in turn influences completed device performance. The parameter space for these performance-influencing factors is large and many-dimensional; this work begins to explore only some of them.

2.4 Other Considerations

Perimeter and Thickness Control

Silicon undergoes an expansion of 10% by volume upon freezing. As samples in this work are sitting flat and subject to hydrostatic pressures when melted, and because the freeze front is moving through the wafer, the end effect is a net mass transfer from the leading edge of the wafer to the trailing edge.

This mass transfer manifests itself in two ways. First, droplets of silicon burst through the trailing edge of the oxide capsule, leaving discrete bumps at the end of the wafer. Second, the thickness of the sample varies: it is thinnest at the leading edge and thickest at the trailing edge. While the former effect is potentially desirable (allowing removal of segregated impurities), the latter is not: it will impact manufacturability of completed solar cells in downstream steps. Recrystallized samples varied in thickness as

much as 100% and no less than 10% (among those measured); this subject will become increasingly important as recrystallized wafers are processed into functioning devices.

This work circumvented dimensional aberrations (particularly in thickness uniformity and surface finish) by mechanically grinding samples down to uniform flatness. However, for commercial viability of this process, dimensional control must be addressed at its root source, not exerted by post-processing. While the exact mechanisms at work here are not immediately clear, they will require further investigation for the commercial success of this project.

References

[1] Brice, J. C. *Crystal Growth Processes*. New York: Wiley, 1986. 29.

[2] Brice, J. C. *Crystal Growth Processes*. New York: Wiley, 1986. 31.

[3] Porter, D. A. and K. E. Easterling. Phase Transformations in Metals and Alloys. Boca Raton: CRC Press, 1992. 15.

[4] Brice, J. C. *Crystal Growth Processes*. New York: Wiley, 1986. 29.

[5] McCormick, James. "Polycrystalline Silicon Technology Requirements for Photovoltaic Applications." *Silicon Processing for Photovoltaics I.* Ed. Khattak and Ravi. Amsterdam: Elsevier, 1985. 13.

[6] Ibid.

[7] Brice, J. C. *Crystal Growth Processes*. New York: Wiley, 1986. 29.

[8] Brice, J. C. *Crystal Growth Processes*. New York: Wiley, 1986. 35-42.

[9] Ibid.

[10] Ibid.

[11] Völkl, Johannes. "Stress in the Cooling Crystal." *Bulk Crystal Growth, Vol 2A*. Ed. D. T. J. Hurle. Amsterdam: Elsevier Science, 1994. 870.

[12] Völkl, Johannes. "Stress in the Cooling Crystal." *Bulk Crystal Growth, Vol 2A*. Ed. D. T. J. Hurle. Amsterdam: Elsevier Science, 1994. 835.

[13] Chalmers, Bruce. "High Speed Growth of Sheet Crystals." *Journal of Crystal Growth* 70 (1984). 3-10.

[14] Roberts, S.G. "Fracture and Brittle-Ductile Transition in Si." *Properties of Crystalline Silicon*. Ed. Hull. London: Institute of Engineering and Technology, 1999.

[15] Ibid.

[16] Siethoff, H. "Macroscopic Mechanical Behavior of Si at High Temperature." *Properties of Crystalline Silicon*. Ed. Hull. London: Institute of Engineering and Technology, 1999.

[17] Chalmers, Bruce. "High Speed Growth of Sheet Crystals." *Journal of Crystal Growth* 70 (1984). 3-10.

[18] Porter, D. A. and K. E. Easterling. *Phase Transformations in Metals and Alloys*. Boca Raton: CRC Press, 1992. 185-190.

[19] O'Mara, Herring, and Hunt. *Handbook of Silicon Semiconductor Technology.* Park Ridge, New Jersey: Noyes Publications, 1990. 4.3.

[20] Porter, D. A. and K. E. Easterling. *Phase Transformations in Metals and Alloys*. Boca Raton: CRC Press, 1992. 274.

Chapter 3: Early-Stage Experiments and Results

Early proof-of-concept experiments varied significantly in scope and process from the final experiments run on thermal profiles and dislocation density. While many early results were qualitative, an overview of these results provides important background for the recrystallization process.

3.1 Capsule and Support Structures

Background

The well-known high surface tension and low viscosity of molten silicon – as well as its wettability of many substrates – indicated the need for encapsulation during recrystallization. Not only would unencapsulated molten silicon tend to ball up due to surface tension, but it would also bond tenaciously to graphite, silicon carbide, or alumina backing plates. Additionally, absent a diffusion barrier, contaminants would be free to dissolve into the melt and damage downstream performance.

There are several important constraints on the capsule: most importantly, it must be pure, low-cost, robust enough to contain molten silicon, and not likely to renucleate the grain structure that preceded it. Given these constraints, the simplest possible capsule is made of a dry thermal oxide, grown directly from the wafer. Thus, early proof-of-concept experiments focused exclusively on SiO_2 as a capsule material. Future work may explore nitride or spun-on capsules.

Uncapped Unbacked Sample without Capsule

Initial samples were 12mm x 24mm, and could sit unsupported between two carrier bars passing through the first-generation recrystallization furnace's hot zone.

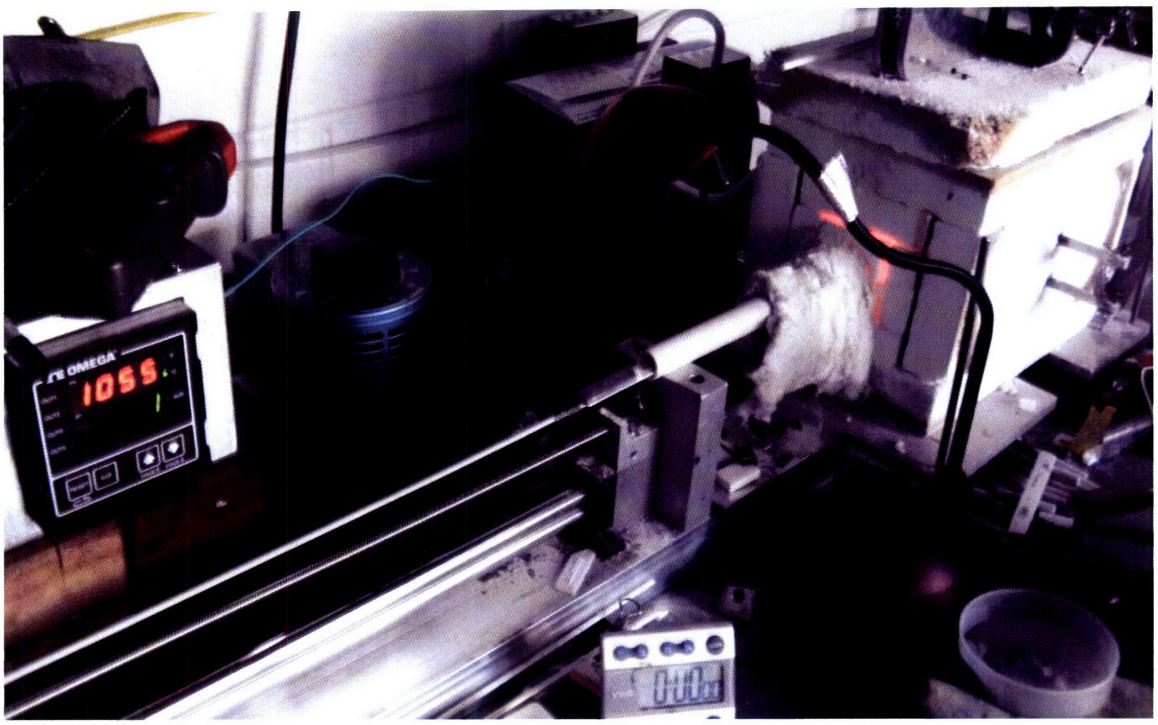

Figure 1: First-revision recrystallization furnace, fitted in this picture with a single carrier bar and backing plate.

To verify the necessity of encapsulating the wafer prior to melting, several samples were recrystallized without backing plates or an oxide layer. The result was extensive capillarity-driven balling. (This early revision of the recrystallization furnace allowed an inert gas purge; samples were recrystallized under argon to prevent any oxide from forming during recrystallization.)

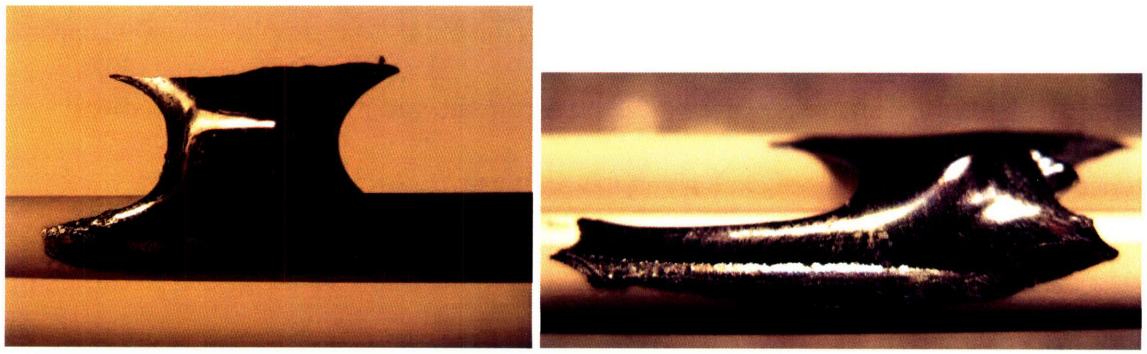

Figure 2: Samples recrystallized without an oxide capsule or backing plates. Recrystallization in these early runs was carried out in an atmosphere purged with inert gas. Starting samples were planar and 300μm thick.

Severe balling and adhesion to the carrier bars was a strong indicator that both a capsule and backing plates would ultimately be necessary.

Uncapped Unbacked Sample with Capsule

The first samples to use oxide capsules employed a two-stage oxidation process: 20 hours of dry oxidation in air at 1100°C, followed by another two hours of steam oxide at 1100°C. Initial tests verified that such an oxide capsule was capable of supporting a molten wafer without leakage.

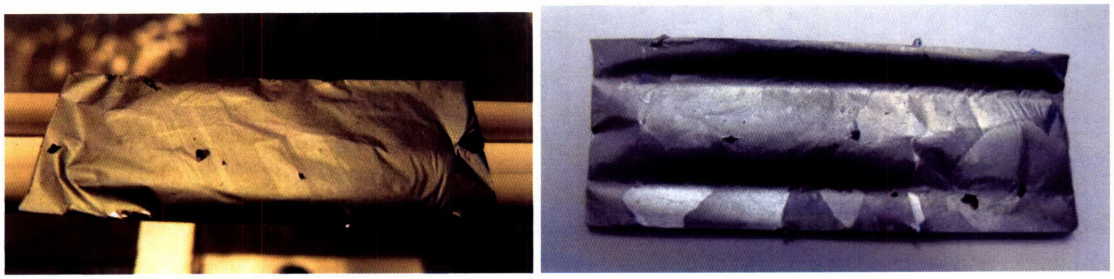

(a) (b)

Figure 3: Top (a) and bottom (b) views of a sample recrystallized without backing plates. The sample has been texture etched in figure (b), to display underlying grain structure. The oxide capsule grown on the sample was sufficiently robust to prevent the molten silicon from spilling, but was not rigid enough to hold the wafer in its planar starting shape. This change in shape is visible in the impression of the carrier bars visible in the bottom of the wafer in (b).

However, the oxide capsule was not sufficiently stiff at melting temperatures to maintain the initial shape of the wafer; it slumped down over the contours of the supporting slider bars. This was a clear indicator that a backing plate on the bottom would be necessary in order to maintain planarity.

Uncapped Backed Sample with Capsule

As an attempt to generate a planar sample without balling, an oxidized sample was melted with a rigid backing plate underneath. The sample remained principally planar, and the capsule helped constrain it from balling. However, the extent of

constraint was insufficient; a meniscus-shaped sample resulted. From this point, the need for an upper capping plate became clear.

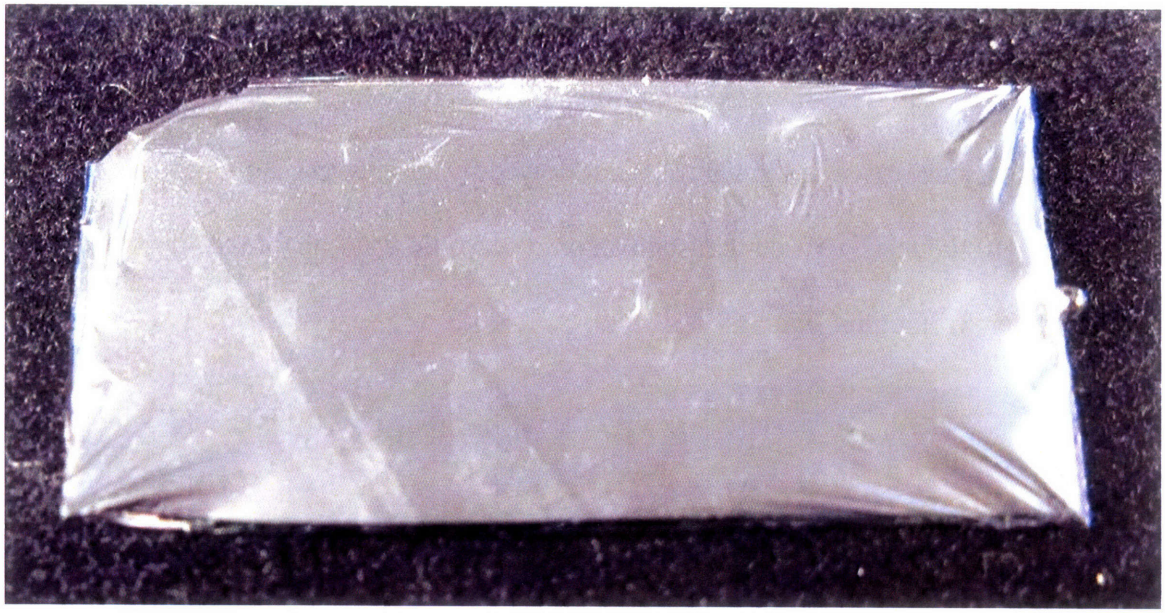

Figure 4: A sample recrystallized with an oxide capsule, on a backing plate, but without a capping plate. Balling is much less severe than without an oxide capsule, but still needs improvement. Thickness variation in this sample was from 200μm to 1000μm.

Backed Capped Sample with Capsule

The configuration that yielded the most consistently successful results was an oxide capsule with both a backing plate and a capping plate. All the samples in the subsequent results were recrystallized in this manner, with two changes from the proof-of-concept experiments.

First, the oxide recipe was modified. Instead of a 20 hour dry plus 2 hour wet oxide, the entire oxidation cycle was reduced to a single 20 hour dry oxidation. Dry oxide tends to be more uniform and of a higher surface quality than wet oxide; as such, its properties are theoretically preferable from the standpoint of preventing new grain nucleation. Second, the furnace was modified such that all samples were recrystallized in

air instead of an inert atmosphere. An oxidizing atmosphere has the benefit of healing cracks in the oxide capsule, preventing leaks and adhesion to backing plates.

Backed Capped Sample Without Capsule in Oxidizing Furnace

Once the recrystallization process was sufficiently robust to run many samples in sequence (for the portion of work outlined in Chapter 4), one further process variation was tested. A sample without a thermally-grown oxide was placed into the recrystallization furnace (at this point, with an air atmosphere). The intention was to investigate whether a separate thermal oxidation step could be avoided, by the rapid thermal oxidation of the wafer as it approaches the melt front in the recrystallization furnace.

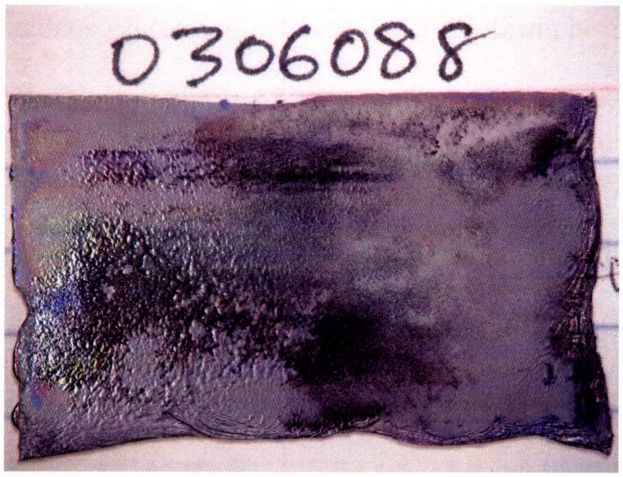

Figure 5: As-recrystallized sample with no pre-grown thermal oxide. Release was successful and no portions of the sample adhered to the backing plates.

The sample released successfully and exhibited similar (if slightly worse) dimensional stability about its perimeter, as compared to samples that were oxidized separately. The ability to avoid a 20 hour oxidation step bodes well for scale-up; long thermal oxidations increase contamination risk, use large amounts of energy, and increase the processing time required per wafer.

3.2 Release Layers

Despite the capsule of SiO_2 surrounding each sample, adhesion to backing plates remained a problem during early experiments. Not only is SiO_2 viscous at the melting point of silicon, but any drops of silicon bursting out the trailing end of the recrystallized wafer (as explained in 2.4) did not have time to oxidize before wetting and adhering to the backing plates.

Several different approaches were attempted to develop a reliable release coating meeting requirements of purity, thickness uniformity, and the ability to firmly register backing plates to a substrate through the release layer.

Boron Nitride

Hexagonal boron nitride (α-BN or white graphite) is capable of high temperatures and is known for its low surface energy. Bare monolithic plates of boron nitride were attempted as an integral backing plate and release layer. Unfortunately, adhesion of the wafer's oxide capsule to the boron nitride was aggressive, to the extent that the wafer shattered upon cooling (due to differential coefficient of thermal expansion (CTE)). Boron nitride was subsequently abandoned as a candidate release material.

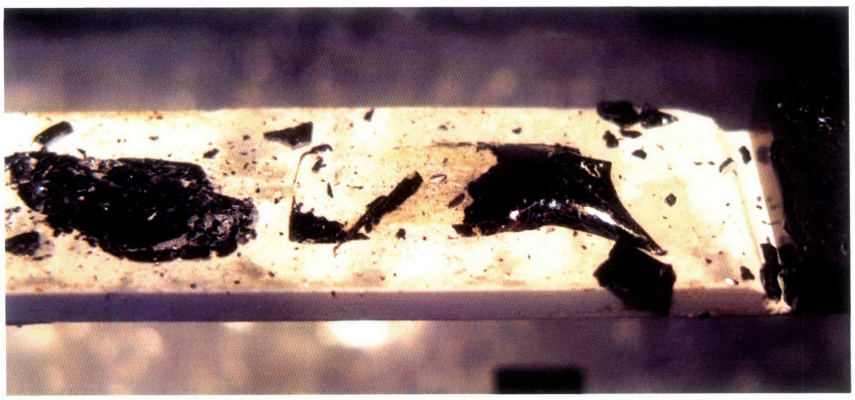

Figure 6: The remains of a wafer recrystallized with monolithic boron nitride used both as a backing plate and release material. The wafer bonded aggressively to the boron nitride and shattered upon cooling to room temperature.

Alumina

The adhesion of the SiO_2 capsule to monolithic Al_2O_3 was also tested, in hopes that a separate release layer might not be required. Results were similarly unsuccessful, with both the sample and the backing material breaking upon cooling. It became increasingly clear that a separately-applied release layer would be necessary.

Figure 7: The remains of a wafer and alumina backing plate, after recrystallization without a separate release layer. Thermal stress from CTE mismatch was sufficient to break both the wafer and backing plate.

Colloidal Silica and Alumina

Spin-coated layers of both colloidal silica and alumina (from Nyacol) were attempted as release layers; the concept was that the thin layer of semi-dense material left by the colloidal solution would shear under thermal stress and allow release of the samples upon cooling. However, both colloidal silica and alumina formulations functioned as strong adhesives rather than release agents. Due to the nanometer-scale particle sizes involved in colloidal suspensions, sintering is rapid. Coupled with the viscous properties of the oxide capsule, adhesion rather than release was promoted. This route was subsequently abandoned.

Graphite

Graphite oxidizes at high temperatures in air, yielding gaseous reaction products. As such, both spin-coated colloidal graphite and thin sheets of solid graphite foil (100 μm thick) were attempted as release layers, with the hope that the boundary layer of exiting gas would prevent the oxide capsule from adhering to the backing plates.

Although combustion was clean and thorough, the entire layer of graphite was visibly gone before the silicon had even reached its melting temperature, meaning that no release layer was present upon melting. Graphite release layers were subsequently abandoned as an area of pursuit.

Slurries

Lacking success with colloidal dispersions of small particles, attempts were made at formulating spin-coatable suspensions of larger ceramic particles. Particles of Si_3N_4, SiO_2, and SiC were all suspended in slurries and attempted as release layers. Constraints on suspension rheology (dictated by the requirement for spin-coating) led to extensive fine-tuning of suspension formulations.

The main vehicle in all suspensions was deionized water. Polyethylene glycol (PEG) was added as a thixotrope. Duramax liquid surfactant was added as a deagglomerant. Octanol was added as a de-foaming agent. Particles were added in various amounts to yield a wide variety of different slurry compositions.

Suspensions were optimized for both spin coating (ability to coat an entire backing plate without gaps) and surface finish. Both of these parameters are functions of viscosity and the surface wetting characteristics of the slurry and the backing plate. Ultimately, to achieve uniform coatings of smooth surface finish, thin release layers were

required; these were achieved with high-viscosity suspensions applied at high (~2000 RPM) spin speeds.

Despite this optimization, the most uniform coatings always led to the most aggressive adhesion. With this knowledge, a pathway of loose powder was explored.

Loose Powder Beds

By spreading a thick bed (400μm) of high-purity dry powder directly onto the backing plate, setting the sample in the powder bed, covering the sample with more powder, and finally setting the upper backing plate onto the stack, the first reliable releases were achieved. All powders attempted (including SiO_2 and SiC of different sizes) succeeded in releasing wafers from the backing plates.

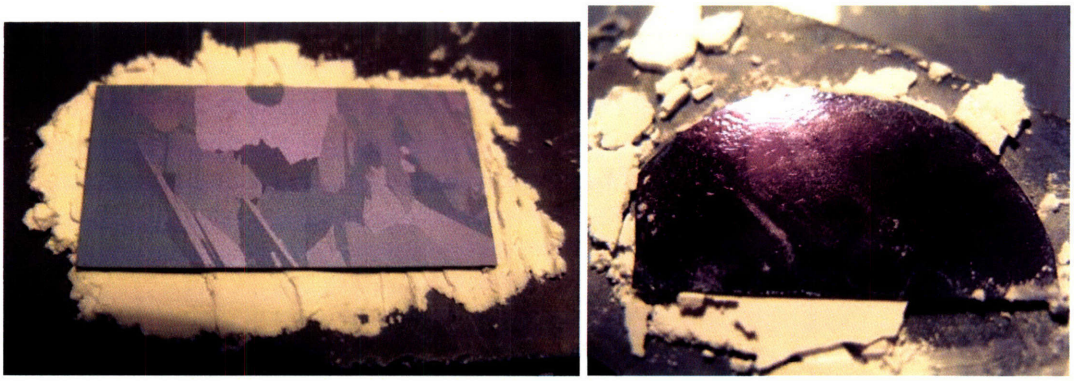

(a) (b)

Figure 8: Unused and used SiO_2 powder beds. The unused bed (a) is spread with a spatula across the backing plate. The used bed (b) is visibly sintered together into solid fragments, which flake off easily for wafer release.

Ultimately, SiO_2 was chosen as a release layer material for two reasons. First, it is transparent to long-wavelength radiation, making radiant heat transfer through the powder bed more readily achievable. Second, the subsequent oxide etch the sample is placed in (to remove the oxide capsule) will also dissolve any remaining particles bonded on from the release layer.

Dimensional Constraint

A major drawback of the thick powder bed method of release was its inability to provide a hard mechanical constraint to the wafer through the powder bed, due to its lack of rigidity. Early samples recrystallized with this method of release exhibited thickness variations of up to 100%. In order to improve thickness uniformity, potential solutions to the problem of mechanical constraint through the powder bed were explored.

Thinning out the beds of applied powder was attempted, but limitations on the precision of manual powder manipulation prevented success in this effort. Similarly, assembling the stack and then nestling the pieces together yielded only thick and discontinuous clumps of silica powder between the backing plates and wafer.

A further attempt was made to clamp down the backing plates, in order to lock in compressive set applied to the powder bed during assembly. Small c-clamps of silicon carbide were placed around a stack of backing plates and spacers to hold the assembly together during recrystallization.

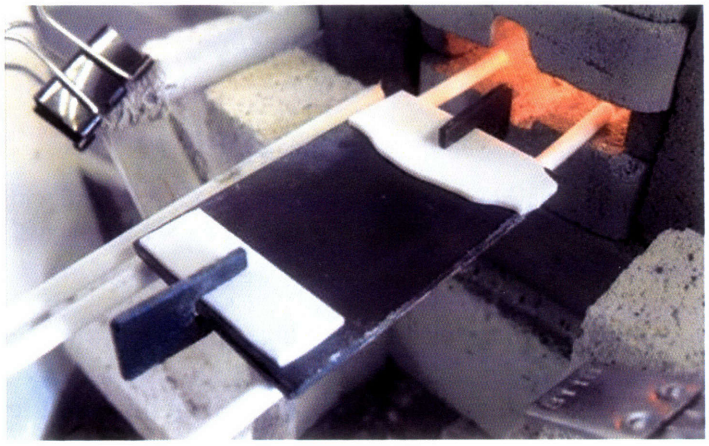

Figure 9: Clamped stack of two backing plates (black), clamps (fork-shaped black pieces), and compressive spacers (white) to ensure rigid fixturing. Underneath the large black backing plate is the sample (not visible in this photo), sandwiched between two compressed layers of silica powder.

Champion thickness uniformity results using this method were worse than champion results using free-floating backing plates; the best documented sample recrystallized between clamped plates was +23% -50% in thickness variation from its starting thickness. The tightest thickness distribution during this set of experiments was carried out between free-floating backing plates (±10%). Free-floating plates required an adhered, uniform powder layer. As such, attention was refocused onto applying thin, adhered, uniform coatings of release powder.

Electrostatic Power Coating

For metal finishing applications, electrostatic spray guns are often used to apply coats of dry polymer powder to metal parts (for subsequent sintering). These spray guns work by entraining particles in a flow of air and imparting an electrostatic charge to the particles on their way out of the gun. By applying an opposite charge to the workpiece, ejected particles will cling to the workpiece surface. Inexpensive versions of these guns are readily available, so electrostatic spray was explored as a means of applying a thin release layer of silica powder.

A generic, low-cost powder coating system was purchased from the distributor Harbor Freight. Initial tests indicated that layers of varying thickness could be easily applied to silicon carbide backing plates and that – furthermore – adhesion of particles was not even dependent on electrostatic biasing. Thin layers and adhesion to backing plates enabled easier stacking (allowing coated plates to be turned upside down without losing powder) and held samples within tighter geometrical tolerances.

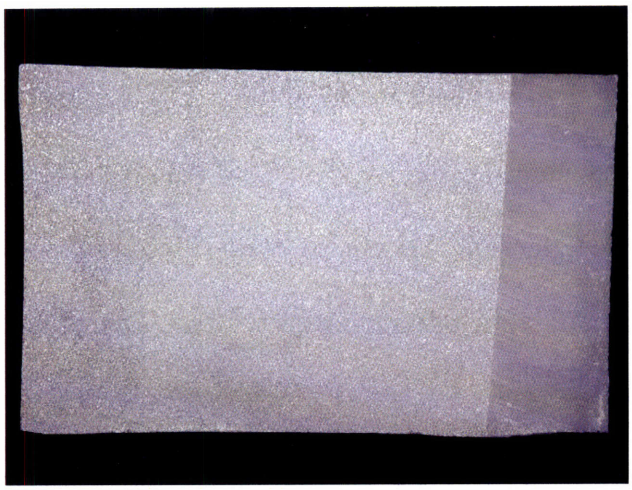

Figure 10: Release layer (light in color) applied to a Hexoloy SiC backing plate via electrostatic spray coating. The darker area on the right was scraped bare with a razorblade to demonstrate the thickness and density of the coating.

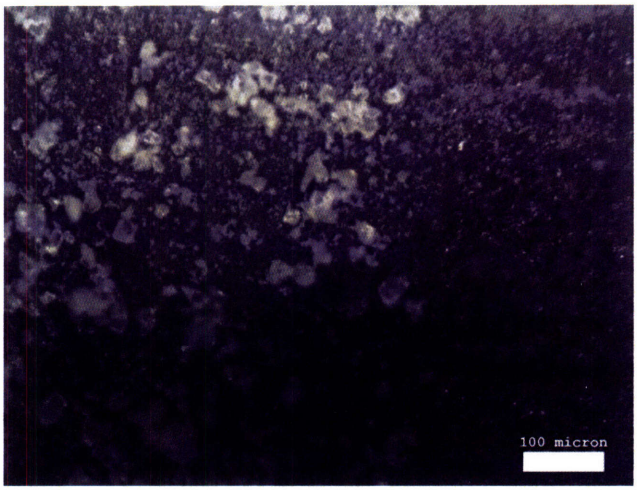

Figure 11: Mircograph of the sample of Figure 10 (unfired), applied by electrostatic spray coating. Although effective as a release layer, clumping and areal non-uniformity are apparent.

For the remainder of proof-of-concept experiments, electrostatic powder coating became the *de facto* method of release layer application, using <20μm silica powder. Subsequent runs employed the technique described below.

Thin-Layer Adhesion and Sintering

Initially, great pains were taken to keep backing plates clean and free of fingerprints. After some time, a backing plate with a fingerprint on it made it into

process; it was found that sprinkled (not spray-coated) silica powder adhered to the fingerprinted region. Furthermore, the coating was sufficiently thin and sufficiently free of clumps (as compared to spray-coated release layers) that it formed a robust, permanent release layer on the backing plate after firing. Area uniformity was also improved over spray coating. Since it provided a smoother surface and the same reliable release, this method was used for the final set of experiments on temperature profile and dislocation density.

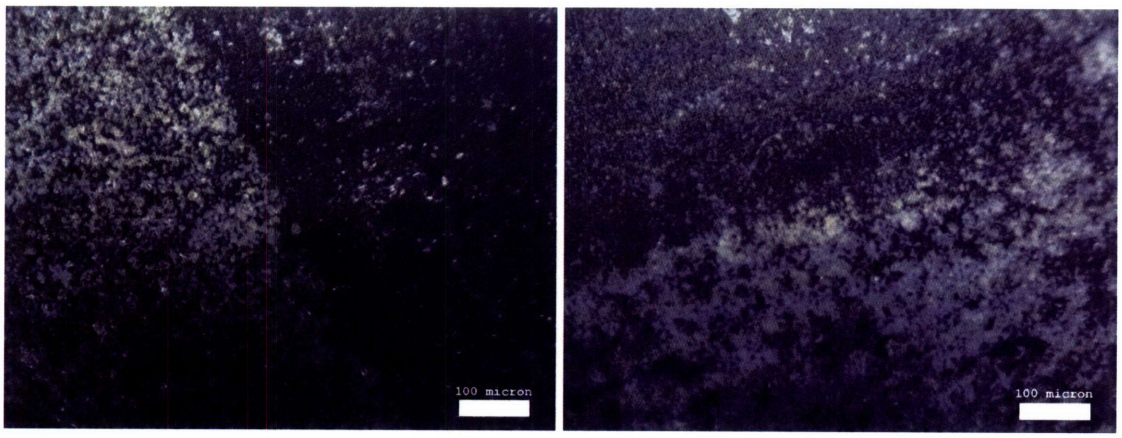

(a) (b)

Figure 12: Micrographs of unfired (a) and fired (b) release layers applied by thin-layer adhesion. Silica particles are white.

Despite its advantages, this technique presents obvious contamination risks; future experiments should seek to replicate the performance of these release layers using a controllable, pure adhesion promoter.

3.3 Backing Plate Material

Fused Quartz

Fused quartz is potentially attractive as a backing material due to its optical transparency (for radiant heat transfer) and purity. However, its temperature resistance is

borderline. Most fused quartz formulations begin to undergo plastic strain above 1200°C. However, they typically do not reach their full rated T_g until 1500°C. Given these constraints, fused quartz was attempted as a backing plate material.

Results were unsuccessful. At the temperatures required for recrystallization, particles from the release layer (at the time, SiC) embedded themselves into the backing plate. Furthermore, the quartz crept to the extent that a secondary backing plate would have been required in order to support the sample to retain its planarity. Finally, adhesion between the wafer and backing plate was aggressive, to the extent that the wafer was destroyed. This route was subsequently abandoned.

Alumina

While alumina is inexpensive and readily available, two issues prevented its adoption as a backing plate material. First, the fact that it contains aluminum was viewed as a contamination risk. Second, the large temperature gradients required in recrystallization imposed large thermal stresses on alumina plates (due to their large coefficient of thermal expansion). After several alumina plates shattered upon traveling through the recrystallization furnace, this route was abandoned.

Mullite

Mullite exhibits superior resistance to thermal shock when compared to alumina, but contamination remains a risk due to mullite's aluminum content. Although several recrystallization runs were successful using mullite backing plates, this route was abandoned due to contamination concerns.

Silicon Nitride

Si_3N_4 is capable of the high temperatures of recrystallization and poses no metallic contamination risks to the process. While several samples were run with silicon nitride backing plates, the lack of a readily-available supplier of thin silicon nitride backing plates led to the abandonment of this material system.

SiC

Like silicon nitride, silicon carbide is capable of temperatures in excess of the melting temperature of silicon and – in certain formulations – contains low levels of metallic contaminants. Saint Gobain's Hexoloy SE is available as a high-purity extruded and ground plate, in a variety of thicknesses and overall dimensions. It is high-density and highly resistant to thermal shock. After repeated success using Hexoloy backing plates, they became the standard backing plate material used in the remainder of this work.

3.4 Grain Structure Replication from Oxide

With release layers and backing plates sufficiently operational, experimental focus was turned to investigating the quality of recrystallized material.

Defining a New Grain Structure

There was initial concern that features from the oxide capsule would nucleate new grains in a pattern similar or identical to the starting grains. To investigate, several experiments were carried out which sought to examine starting and finishing grain structure. In the first set of experiments, cast multicrystalline silicon was recrystallized, and before-and-after images were compared. Results indicated that – indeed – a completely new grain structure resulted from recrystallization.

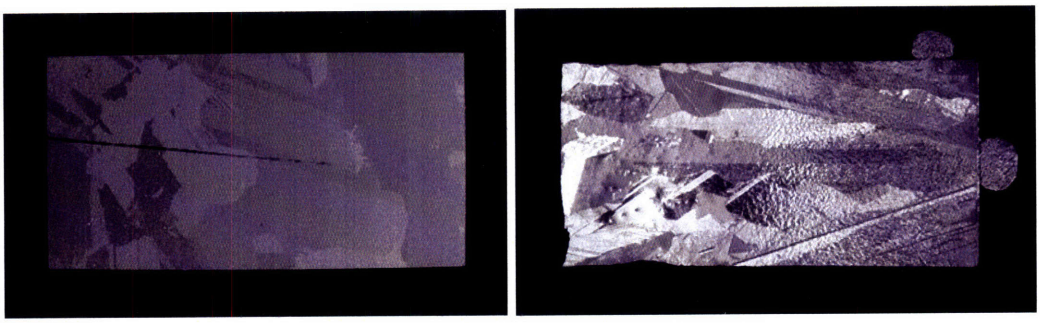

(a) (b)

Figure 13: Starting wafer (a) and the same wafer, recrystallized and texture-etched (b). Definition of a completely new grain structure is apparent.

In the second set of experiments, a 50mm diameter monocrystalline wafer was recrystallized. In the same way, new grains formed, indicating that the oxide capsule does not dictate the grain structure of the recrystallized sample.

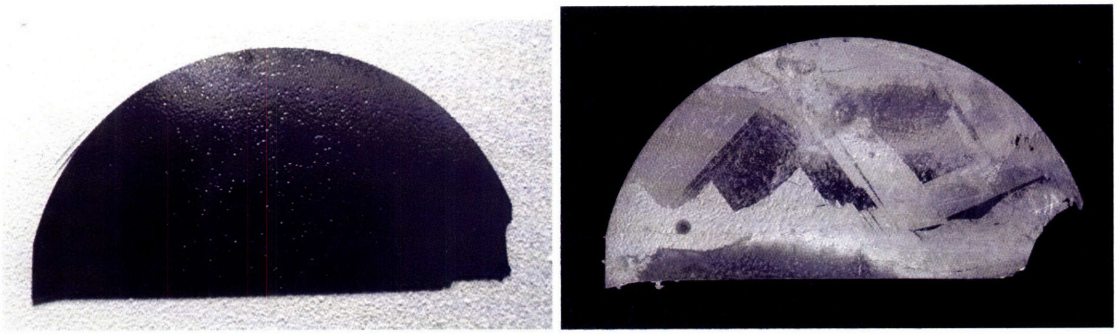

(a) (b)

Figure 14: Monocrystalline wafer before (a) and after (b) recrystallization and texture etching. The after picture (b) clearly shows the definition of new grains.

Top-Bottom Grain Structure Agreement

As explained in 2.3, in-plane grain boundaries are ideally avoided. Several runs sought to verify agreement in grain structure between the top and bottom of recrystallized wafers; this verification would bode well for the prospects of actually improving electrical performance.

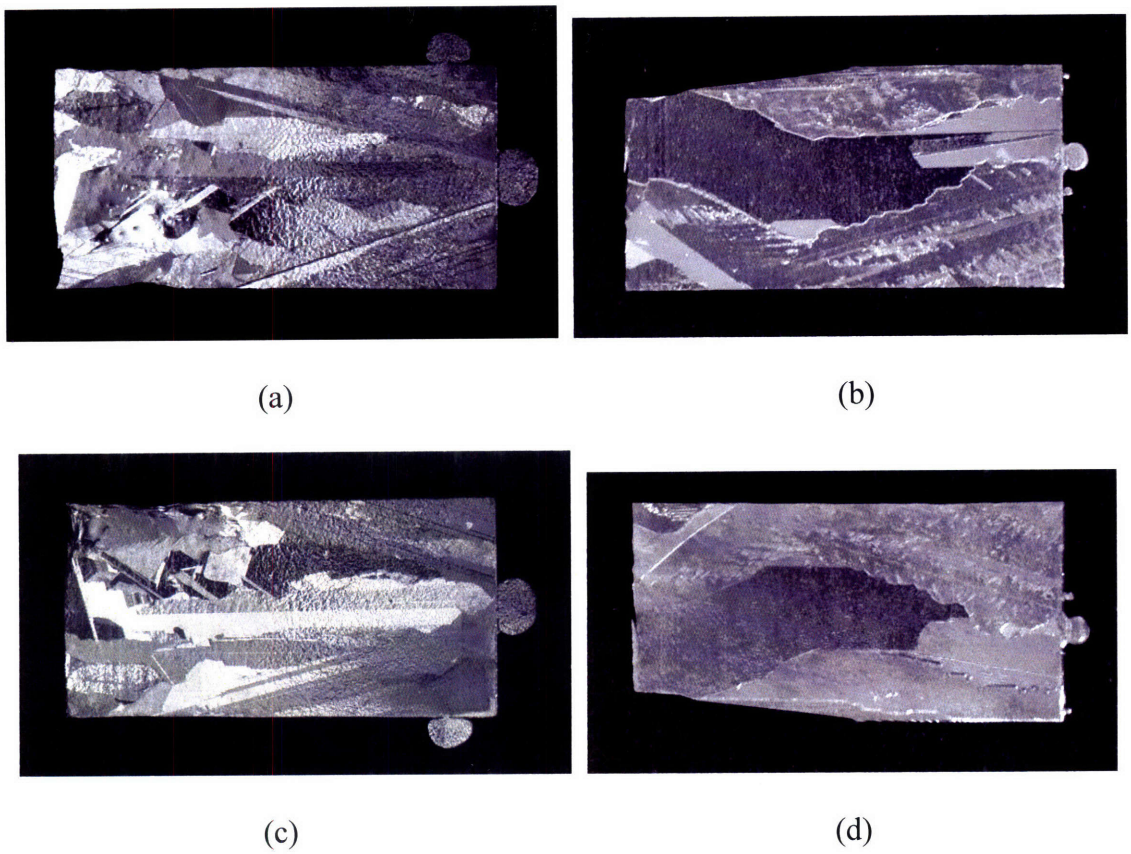

Figure 15: Two recrystallized multicrystalline samples. (a) and (c) shows the top and bottom of one sample, (b) and (d) show the top and bottom of a different sample. Samples were rotated about their long axis from front to back. The sample of images (a) and (c) was recrystallized between clamped backing plates with a spray-coated release layer. The sample of images (b) and (d) was recrystallized between clamped plates in a thick powder bed. Similar grain shapes and locations appear top-bottom in each sample, which indicates that most grain boundaries are not parallel to the plane of the wafer. Temperature control was open-loop at the time of this recrystallization, which accounts for the difference in grain structure between samples (a) and (b).

Preservation of Surface Features

In its ultimate commercial application, this process may be applied to wafers bearing predefined structures, such as light-trapping grooves. Thus, there is interest in recrystallizing samples which include microscopic surface features and investigating the preservation of geometry.

To investigate this, a monocrystalline wafer with etched 20μm light-trapping grooves was recrystallized.

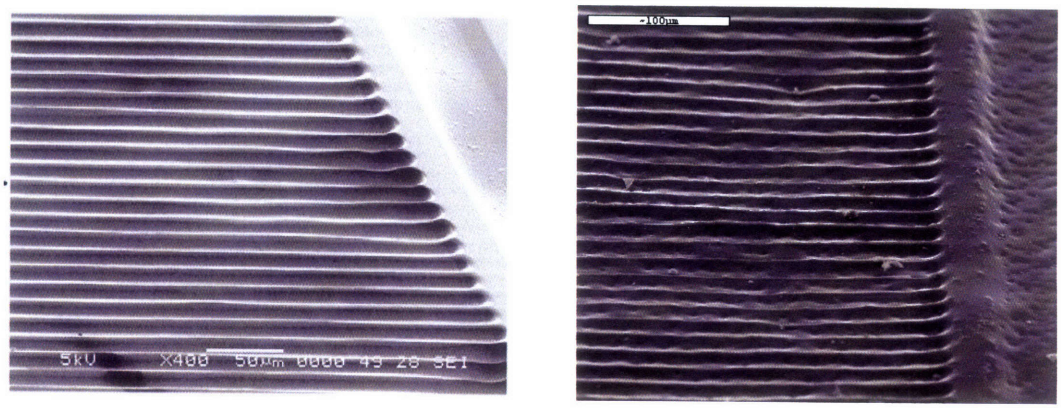

(a) (b)

Figure 16: SEMs of a textured sample, before (a) and after (b) recrystallization. Although some features were damaged in recrystallization, the overall shape was preserved.

Results indicated that surface features can be largely preserved during recrystallization. This result is a promising starting point for further investigation into geometric fidelity in recrystallization.

Chapter 4: Hardware

Unique constraints of the recrystallization process necessitated new hardware development. Few off-the-shelf furnaces reach the temperature ranges required for the melting of silicon (in excess of 1414°C); those which do usually require an inert gas or reducing atmosphere. In line with managing the economics of the recrystallization process, it was sought to eliminate the use of inert or reducing gases from the process flow – rendering these off-the-shelf furnaces inapplicable to this work.

Furthermore, owing to the unique purity requirements of PV applications, the metal and metal oxide components present in much off-the-shelf furnace hardware rendered it useless for recrystallization, particularly given the temperature ranges in use. In light of both of these factors, it was necessary to design and build a custom furnace for recrystallization processing.

4.1 Elements and Contacts

Material Constraints

The high concentration of metallic impurities (such as aluminum, iron, and nickel) in off-the-shelf SiC heater elements require that unconventional SiC heater elements be used in the recrystallization furnace. Specifically, Saint Gobain's Hexoloy SE was chosen as a heater element material for its low levels of metallic impurities. While a competing material -- Poco's SuperSiC -- possesses similarly low levels of contaminants, the smallest diameter readily available (1/4") was too large to accommodate the spacing and quantity of heater elements desired in the furnace. Thus,

Hexoloy was the best choice at the time. In the future, custom shaped SuperSiC elements should be employed to eliminate the cold-start issues associated with Hexoloy.

Conventional SiC heater elements, such as IsquaredR's StarBar, employ infused silicon and metals to control resistivity along the length of the heater element. For instance, by coating the ends of the element with aluminum, StarBar elements run cool at the ends and hot in the center. As Hexoloy SE is not intended by the manufacturer as a heater element material, no such contacting provision is made on stock material. The entire heater element runs hot, out to the extreme ends. As a result, much work went into making reliable electrical contact to the elements: the temperature was too high for graphite in air, and refractory metals such as tungsten and molybdenum pose a large transition metal contamination risk.

Contacts: Metal and Graphite in Air

Initial attempts at electrical contact involved clamping steel and graphite plates with large surface area to the ends of the heater elements. It was intended that enough heat would be radiated and convected out of the large-area contact plates to keep the metal-SiC interface cool. This would either allow the use of graphite as a contact material, or keep the metal cool enough to eliminate concerns about contaminant diffusion into the heater element.

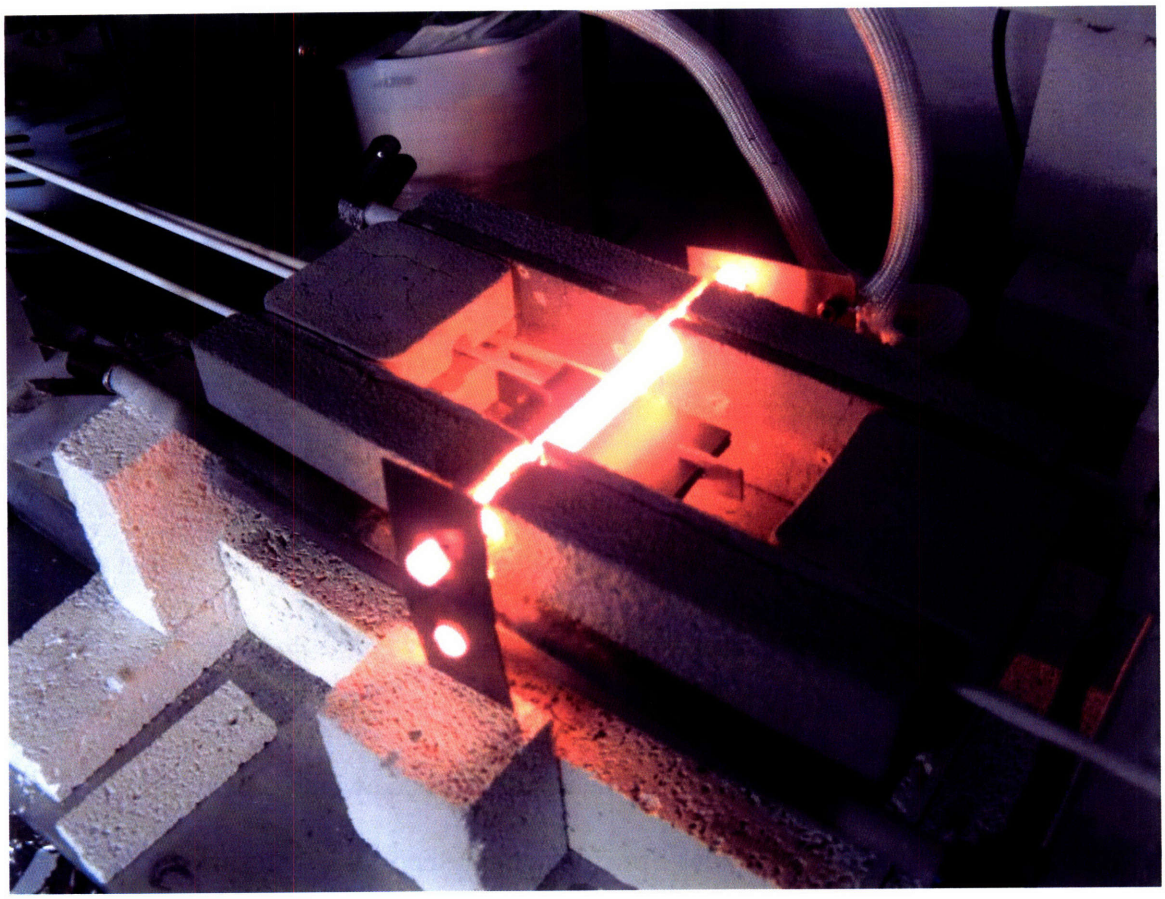

Figure 1: An early furnace using interference-fit metal plates as electrical contacts. Extensive local heating around the contacts is apparent from the element's bright glow; this behavior causes rapid contact failure.

While these early iterations were successful in short runs, two large drawbacks arose over the long term. First, the large heat flux out of the ends of the elements required that large amounts of power were dissipated, in order to achieve sufficient temperature at the middle of the element. Running smaller contacts with less parasitic heat loss only caused the contacts to oxidize and fail quickly. Second, the steep thermal gradient from the end plate to the center of the element placed large thermal stresses on the SiC. Over time and sufficient thermal cycles, the SiC heater element broke near the contact interface.

Silicon Carbide Contacts

From this point, SiC contact material was explored, owing to its higher temperature capabilities. Early tests using sawn-off, aluminum-coated ends of StarBar heater element were successful (but unreliable), and they required a physically bulky clamp mechanism in order to ensure contact throughout the temperature range in question. In addition, electrical contact force varied significantly between room temperature and operating temperatures – decreasing with increasing temperature – due to differential thermal expansion.

In order to deal more effectively with temperature-dependent contact resistance and contact degradation, a spring-loaded axial design was attempted. In the same manner as a carbon arc lamp's spring-loaded contacts provide consistent electrical contact to electrodes, an axial spring-loaded arrangement in the furnace would provide reliable contact and a mechanical bias toward firm contact. Furthermore, with sufficient spring travel built in, the gap caused by element erosion at contact points could be picked up by travel of the contacts.

Early tests using this arrangement were promising, but several shortcomings derailed its implementation into the final furnace design. First, oxidation at the interface between the large-diameter contact bars and the small-diameter heater bar was so severe that startup was impossible without abrading the contact surfaces beforehand. In addition, even small amounts of axial load were sufficient to induce buckling in the heater elements while at temperature. Over time, elements became severely curved. This approach was subsequently abandoned.

Element Shape

Cool heater element ends can be achieved by varying the cross-sectional area of the element, i.e., necking down in the center. While creating a monolithic element would have eliminated the need for joining disparate pieces of SiC together, lead times and costs precluded ordering custom Hexoloy or SuperSiC elements. Instead, efforts were made to contact large-diameter SiC to small-diameter SiC.

Early attempts used lugs made of Poco's SuperSiC, and attempted to braze these lugs onto the ends of a Hexoloy bar. Pulverized P++ Si (.005 Ω-cm) was used as a brazing alloy. In both polyethylene glycol-based paste and powder form, attempts at brazing were generally successful in mechanically fastening two pieces of SiC together, but the joint was not strong (breakable by hand). In addition, the brazing had to be carried out under a reducing atmosphere (forming gas) and with thoroughly oxide-etched Si particles (which required extensive washing). Finally, the electrical resistance of brazed joints was high. This route was subsequently abandoned.

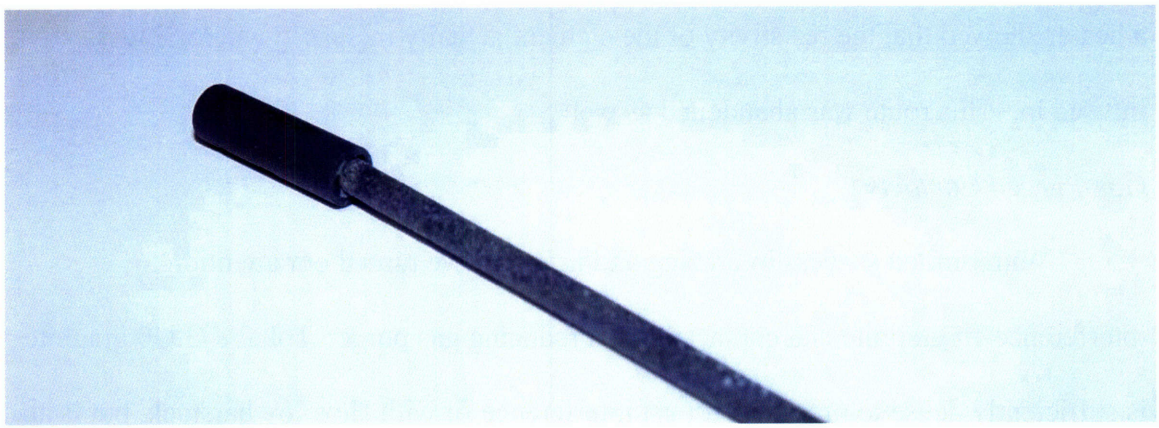

Figure 2: Brazed end lug. The large diameter section is of low-resistivity SuperSiC, while the small-diameter section is high-resistivity Hexoloy SE. Silicon was used as a brazing material.

Variable Resistivity

A few experiments sought to replicate the StarBar's variable resistivity by selectively infusing Si into the porosity of a pure SiC heater element (in this case,

SuperSiC due to its higher porosity). In theory, infusing silicon into the bulk of a porous SiC element should decrease the resistivity locally. A graphite crucible held pulverized P++ Si in contact with the midspan of a length of SuperSiC.

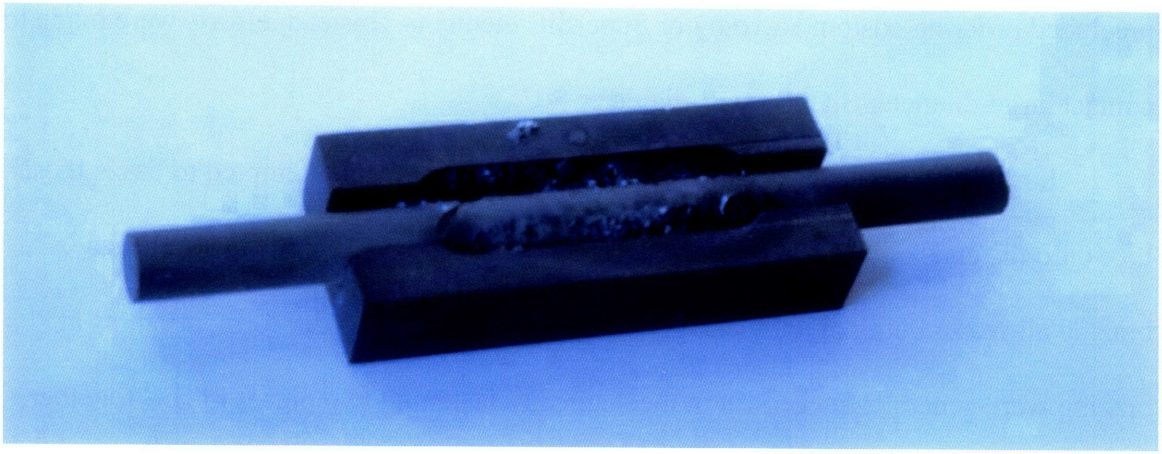

Figure 3: Attempt at infusing silicon into SiC in order to change its resistivity. The rectangular graphite block was filled with silicon granules, while the SiC was laid down through the center, in contact with the granular silicon.

Firing was again carried out in a tube furnace under inert gas, with oxide-etched particles. While wetting and infusion were apparently successful, running the element as a heater showed that the resistivity of the element actually increased where Si was infused in. This route was abandoned as well.

Gas-Purged Contacts

With limited success in brazing and infusing, we turned our attention to interference-fit graphite end contacts with a reducing gas purge. Tokai's G530 graphite is sufficiently dense to provide a robust interference fit with Hexoloy barstock, but is also sufficiently porous to allow purge gas to flow through its bulk. By drilling out the end contacts and purging them with forming gas, oxidation was slowed to the extent that graphite functioned as an effective contact material. A shielding tube of Hexoloy was added around the graphite lugs, in order to channel purge gas more effectively through

the graphite. Finally, the length of each graphite lug was increased, allowing enough outward heat transfer for exposed graphite electrical contacts at the ends of each element. The graphite lugs which protrude from the ends of each element are contacted by Tokai G530 graphite split clamps, machined to fit around the lugs. Nickel ring terminals are bolted to the split clamps, which provide electrical contact to the busbars through high-temperature wires.

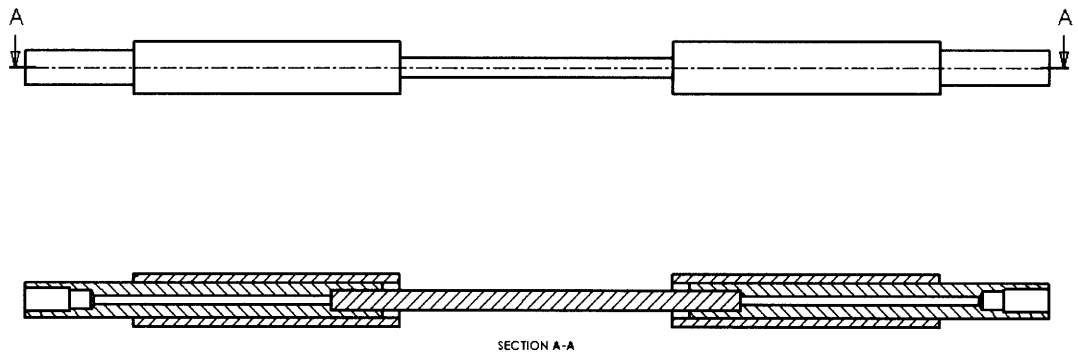

Figure 4: Planview and cross-sectional view of heater element design employed in the recrystallization furnace. The negative-sloped hatch marks represent graphite; positive-sloped hatch marks represent Hexoloy SE silicon carbide.

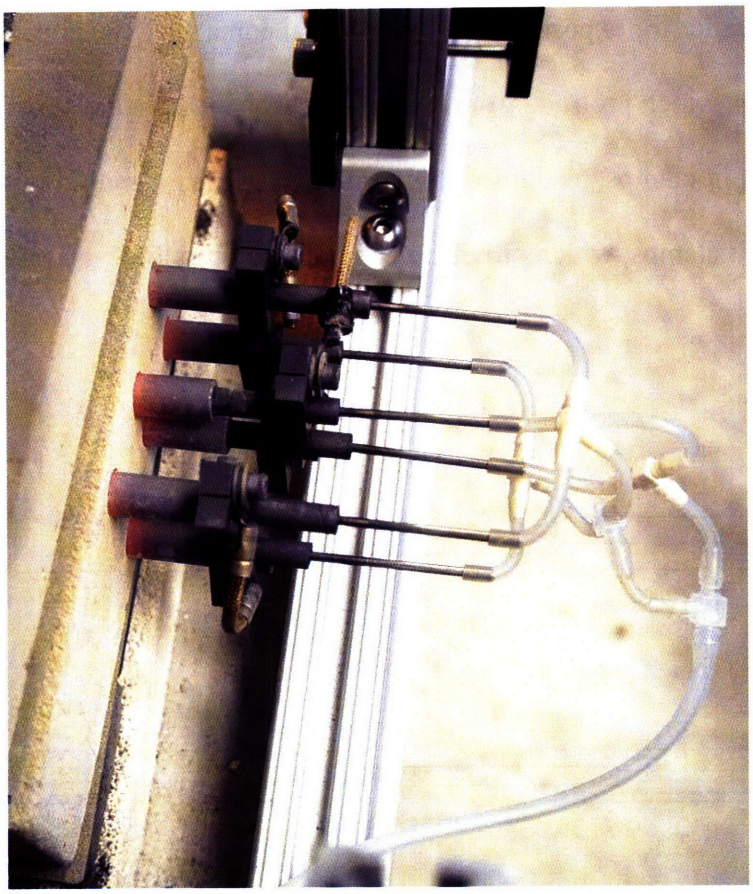

Figure 5: Top view of the element contact design on one side of the furnace. Purge gas flows through the translucent silicone tubing and into the graphite lugs through stainless barbs. Because of its spacing from the hot portion of the element, the exposed graphite components are able to last for tens of hours in air.

The current design employs 1.4 bar (at the regulator) forming gas purge in all end lugs, at a flowrate of 5 standard liters per hour (SLH) per end lug. Despite the purge of reducing gas, element life remains short – on the order of 5 hours at operating temperature. Failure typically occurs at the interference-fit interface between graphite and Hexoloy. Future designs should employ custom-made monolithic SuperSiC elements in order to eliminate SiC-graphite interfaces and the need for purge gas.

A further drawback of the current design is that the Hexoloy SE barstock employed as a heater element begins to decompose into SiO2 and CO2 around 1650°C (which contradicts the manufacturer's rated 1900°C capability, in air).

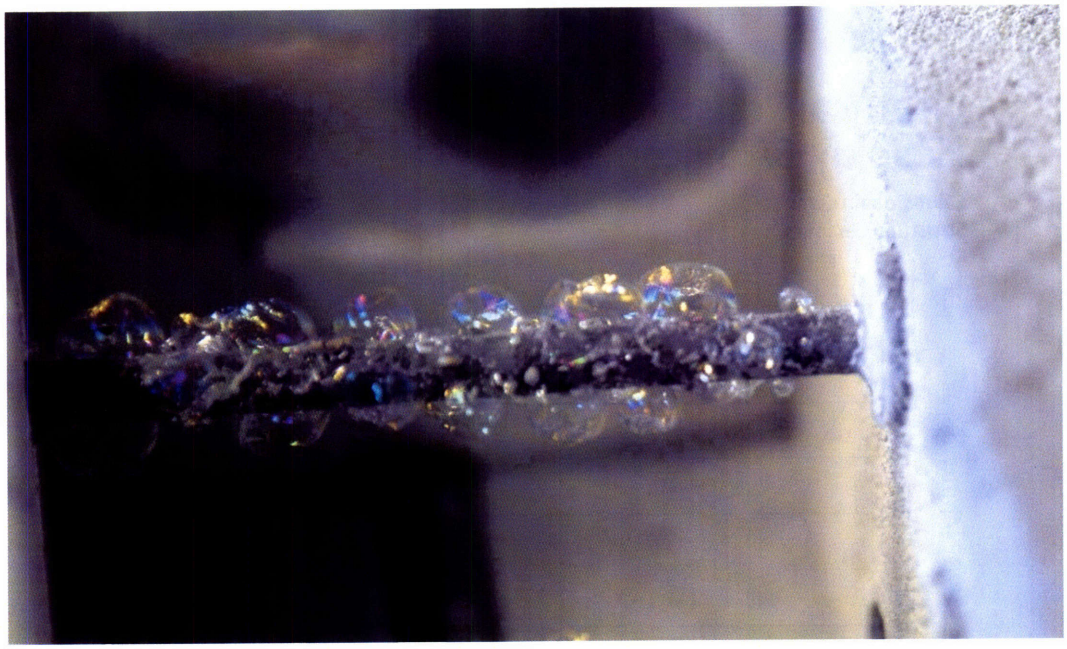

Figure 6: SiO2 bubbles formed from element decomposition.

Although these bubbles do not influence the accuracy of the optical temperature measurement used for control (as ratio pyrometers are used), their formation indicates rapid element degradation. Further iterations should seek either to use a different formulation of silicon carbide, or to go to a larger diameter heated portion. Increasing the diameter of the heated length would increase the radiant area for heat transfer, thereby decreasing the required element temperature and prolonging element life.

4.2 Control

Hexoloy SE's resistivity is not linear with temperature; the resistivity is sufficiently large at room temperature that safe voltages (240V and below) are insufficient to light it off. As temperature increases, resistivity falls off steeply,

particularly above 1600°C. These constraints have required several modifications: control off a pyrometer as opposed to a thermocouple, control of current instead of voltage, room-temperature startup load, dedicated independent loops, and a kicker heater to jump-start the furnace.

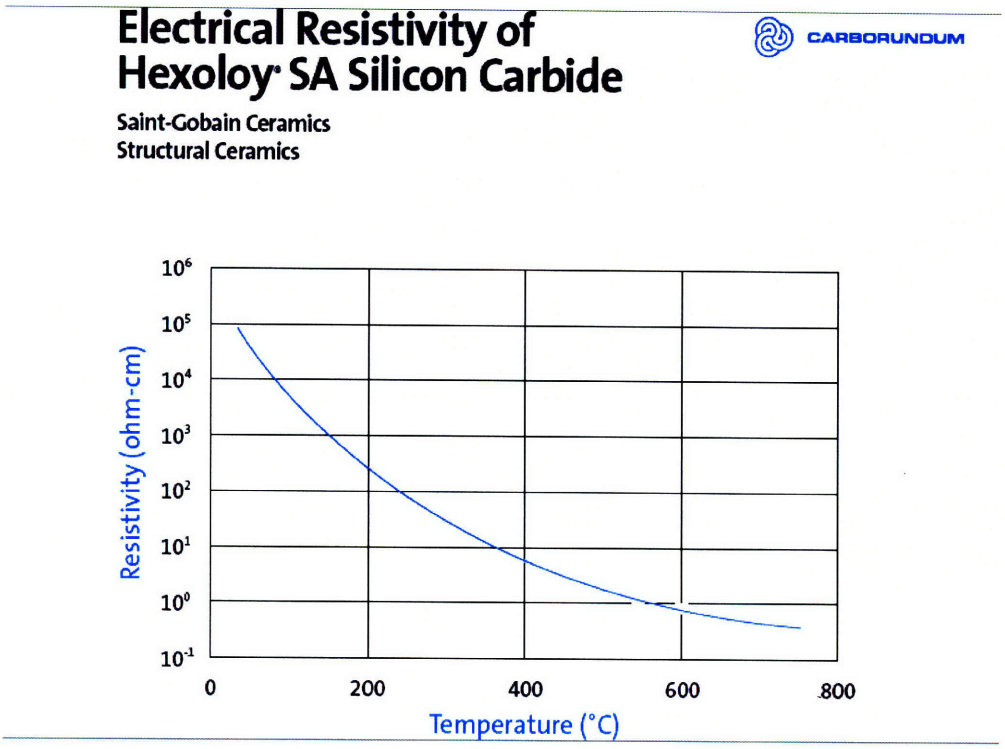

Figure 7: Resistivity of Hexoloy SA (chemically identical to Hexoloy SE) with temperature. The rapid falloff toward 0Ω above 800°C placed unique constraints on furnace electrical design.

Configuration

Elements are configured in 3 zones. Each zone is driven off a 240VAC rail, throttled by a phase-angle fired SCR power controller (Control Concepts, Inc.'s 1032 series 40-Amp controller). Early furnace runs suffered from zone current saturation upon startup, leading to overheating and early element failure. To correct this problem, an 8A current limit was implemented on each of the power controllers. Controller span was adjusted accordingly. While the controllers (by their design) still saturate upon startup,

the 8A limit protects against overheating while still allowing enough power to handle thermal transients.

Each zone consists of 2 elements in series, for a total of 6 elements. Series pairs are arranged vertically from one another, such that independent zones occur along the travel direction of the sample moving through the furnace. Due to the series connection in each zone, top-bottom thermal symmetry in the furnace requires equal wear on the top and bottom elements (as element resistivity tends to increase with hours of operation). As such, a concerted effort was made to pair elements with similar numbers of hours of operation.

Furthermore, as elements of different ages possess different resistances, it was not possible to control all elements from a single loop. Older elements run hotter than newer elements carrying the same current, so zone-to-zone temperatures will vary widely if controlled from a single loop. Consequently, three independent loops were used.

Startup

Hexoloy's resistivity at room temperature is on the order of 10^5 Ω-cm. Resistivity decreases with increasing temperature, but the amount of current driven from a 240V rail at room temperature does not generate sufficient self-heating to sustain a temperature ramp. As such, a kicker heater element was incorporated into the furnace in order to bring elements up to ramp-capable temperature from a cold start. The kicker element – a length of Poco's SuperSiC silicon carbide -- is recessed into the floor of the furnace. On startup, it is powered on and preheats the furnace temperature to approximately 750°C, at which temperature the main heater elements can begin to conduct and take over. At this point, the kicker element is powered off.

A further disadvantage of Hexoloy's high cold-start resistivity has to do with the SCR power controllers: they are unable to sustain any power output without at least 250mA of load. As the SCR effectively sees an open circuit at startup (no load current through the elements), it was necessary to implement parallel shunt loads for startup. A 240V incandescent light bulb (effectively an inexpensive and safe power resistor) across each zone pulled enough load current for the SCR to successfully light an element.

Owing to the positive temperature coefficient α of tungsten light bulb filaments, it is necessary to disconnect the light bulbs once the elements have begun to conduct. While the negative α of Hexoloy should in theory work counter to the light bulb's positive α, the scale of these responses is vastly different and the control loop is unable to stabilize with the light bulbs connected.

Pyrometers vs. Thermocouples

Owing to the negative α of Hexoloy SE, elements under full rail voltage can swing from 250mA to 8A of current in less than 4 seconds. While thermocouples are an understood and reliable method of feeding back temperature for control, their thermal mass (particularly when coupled to a ceramic thermocouple tube) is too large to accurately track rapid swings in element temperature. Furthermore, thermocouples are unable to measure actual element temperature: a thermal short circuit down the ceramic thermocouple tube drags down thermocouple bead temperature, rendering the measurement inaccurate. Finally, the type R thermocouple necessary to measure temperatures up to 2000°C contains platinum and rhodium, which are damaging contaminants for silicon-based semiconductors.

In light of these considerations, pyrometers were implemented as a means of temperature feedback for furnace control. The pyrometers were of a custom design, employing two detectors tuned to different wavelengths and an algorithm to calculate temperature independent of emissivity. As opposed to single-color pyrometers (which infer temperature based solely on the intensity of incoming radiation), the two-wavelength design employed in this furnace instead senses the "color" of incoming light, rendering measurements impervious to emissivity changes and optical path blockages.

The pyrometers employed collimating light traps made of graphite, in order to reject off-axis stray radiation. As installed, these light traps ensured that a spot size of 3mm was sampled directly off the heater element, preventing view of other surfaces inside the furnace.

Incident radiation from the heater elements is split by a thin piece of microscope cover glass: approximately 8% of the light is reflected into the long-wavelength (IR) detector, while the remaining 92% is transmitted to the short-wavelength (visible) detector. On-board transimpedance amplifiers output a voltage proportional to the short-circuit photocurrent out of each detector. These signals are fed to a board containing scaling and processing circuitry.

The algorithm was implemented on board a Microchip DSPIC30F2012 microcontroller, programmed in C. In each loop, analog voltages from both color channels were prescaled with non-inverting op-amp circuits, converted to digital using Microchip's on-board 10-bit A-D converter, and then fed back out on a USART serial bus to an Analog Devices 10-bit D-A converter, to represent a linearly scaled voltage (0-5V) corresponding to temperature. This 0-5V signal feeds into the Omega CN3251

temperature controller, which in turn drives the power controller, providing current to the element. Analog voltages from the pyrometer are sampled at 20 Hz, providing rapid temperature feedback.

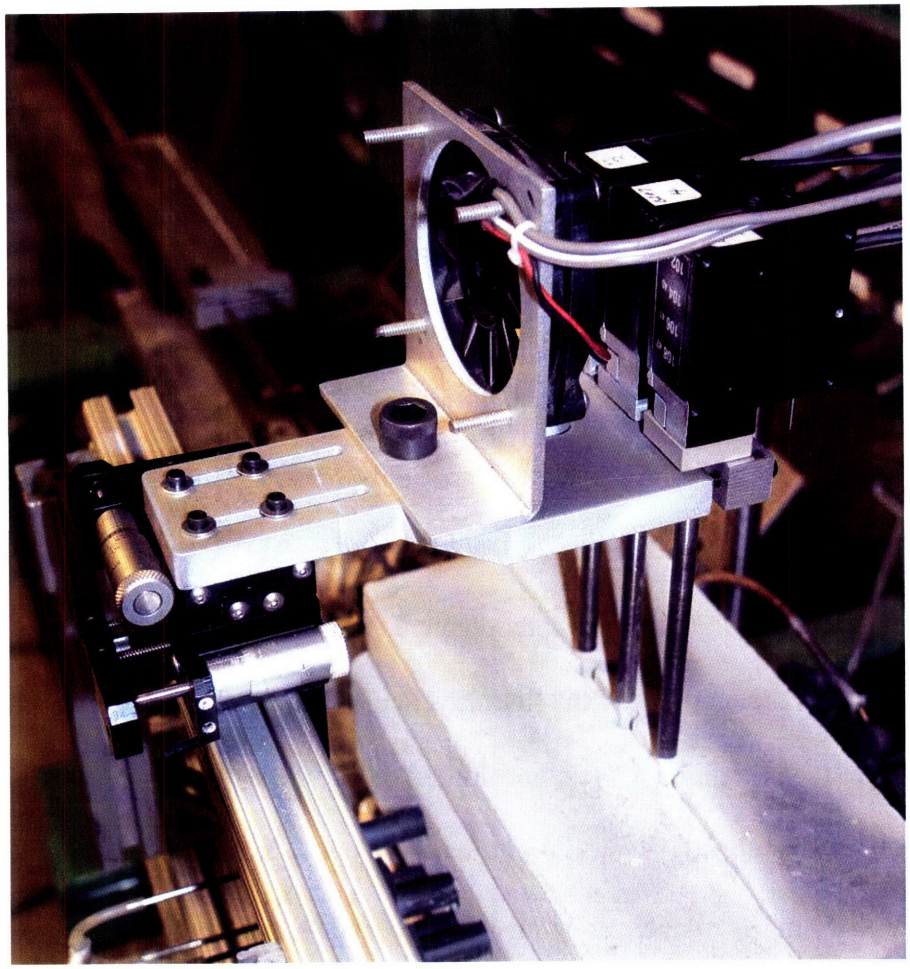

Figure 8: Pyrometers as mounted on the furnace. The collimators are aimed through holes in the roof of the furnace and look directly at the heater elements. An XY micrometer stage allows for adjustment of aim. A fan is used for forced convection cooling of the pyrometer housings.

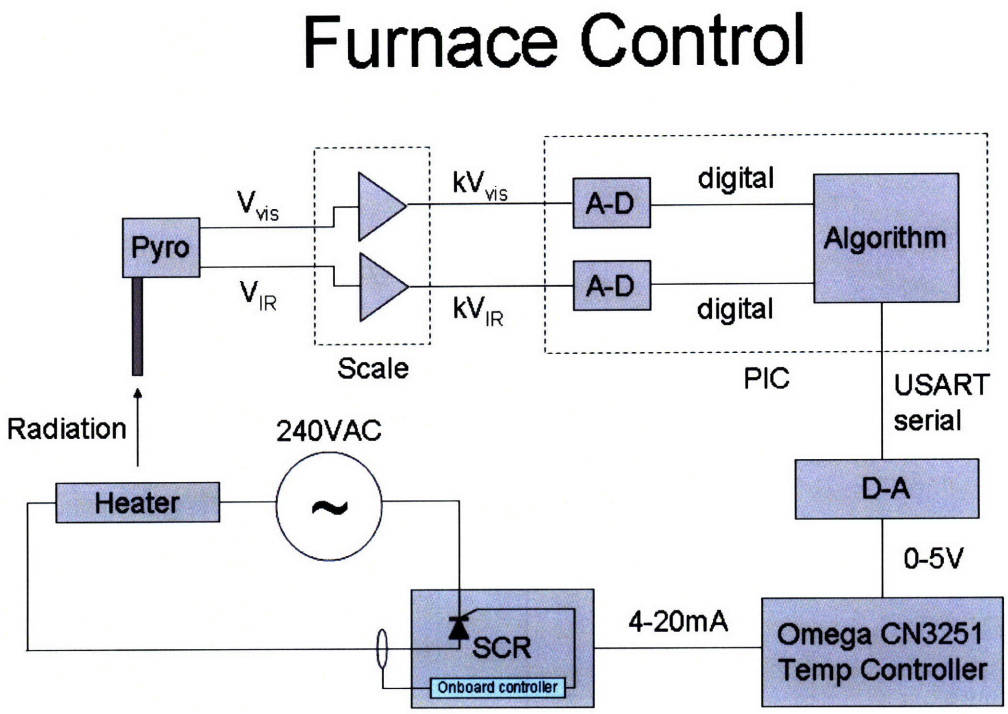

Figure 9: Schematic of a single control loop in the furnace. There are 3 such loops.

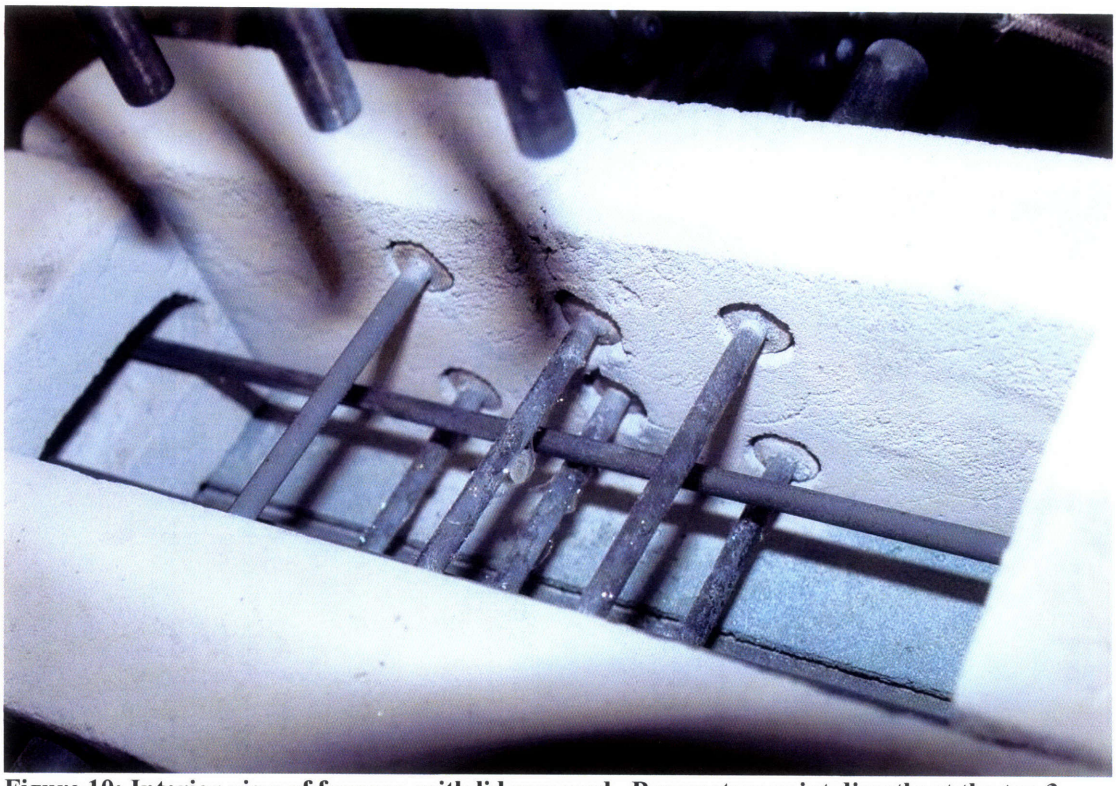

Figure 10: Interior view of furnace, with lid removed. Pyrometers point directly at the top 3 elements. The long bar running left to right is one of the sample slider bars.

Current vs. Voltage control

Power dissipation in a resistive heating element is

$$P = I^2 R = \frac{V^2}{R}.$$

Control Concepts, Inc.'s 1032 SCR power controller is operable in standard voltage control mode as well as current control mode (if the main board is replaced). In voltage control mode, an increase in voltage dV exerted by the power controller increases total power by more than the factor $\frac{(V+dV)^2}{V}$ because the temperature coefficient α of the heater element is negative and R is in the denominator of the power equation. (The opposite case is true with standard nichrome and tungsten heater elements, which are more common; α is positive for these materials.) Consequently, the system is prone to saturation and overshoot in voltage control mode, particularly on startup.

In current control mode, an increase in loop current dI exerted by the power controller is accompanied by a power increase of *less* than the factor $\frac{(I+dI)^2}{I}$; the accompanying decrease in R of the heater element decreases the effect of a control action. As a result of this behavior, it was necessary to configure the loops to operate in current control mode, as stable operation in voltage control mode – while possible – was unreliable.

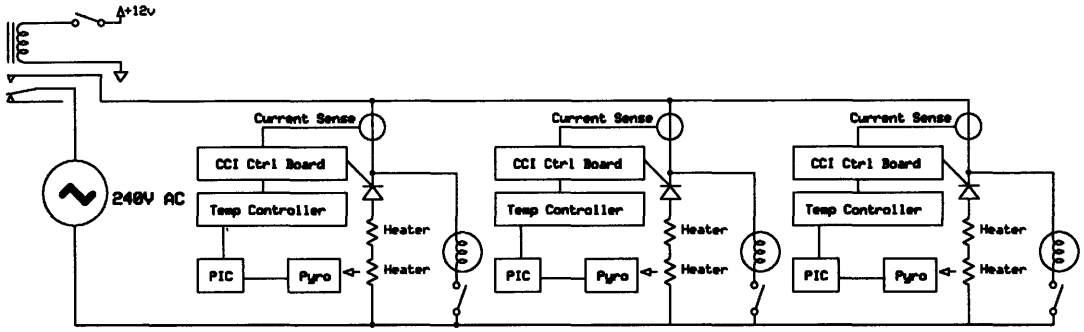

Figure 11: Schematic of the final furnace circuit, as built. The light bulbs in parallel with each pair of heaters were turned during startup to allow current draw out of the SCR. The CCI control board, current sensor, and SCR are all integrated into the CCI 1032 power controller package.

Control Loop Parameters

The Omega CN3251 temperature controllers used to regulate element temperature provide full PD control along with integral reset time adjustment. Integral gain is fixed.

While the controllers are potentially fast enough to control the system (2 Hz sampling rate), they struggled to keep pace sufficiently with rapid transients. As a result, a highly conservative control scheme was implemented, with a wide proportional band (625°C), no derivative control, and a brief integral reset period (6 sec., to minimize rollup). The resulting control was successful in tracking temperature setpoints to within 2°C steady-state. While underdamped in responding to transients, sample motion through the furnace was sufficiently slow that the furnace had little trouble maintaining

setpoints. Indeed, the largest transients the furnace is subject to are on startup and shutdown.

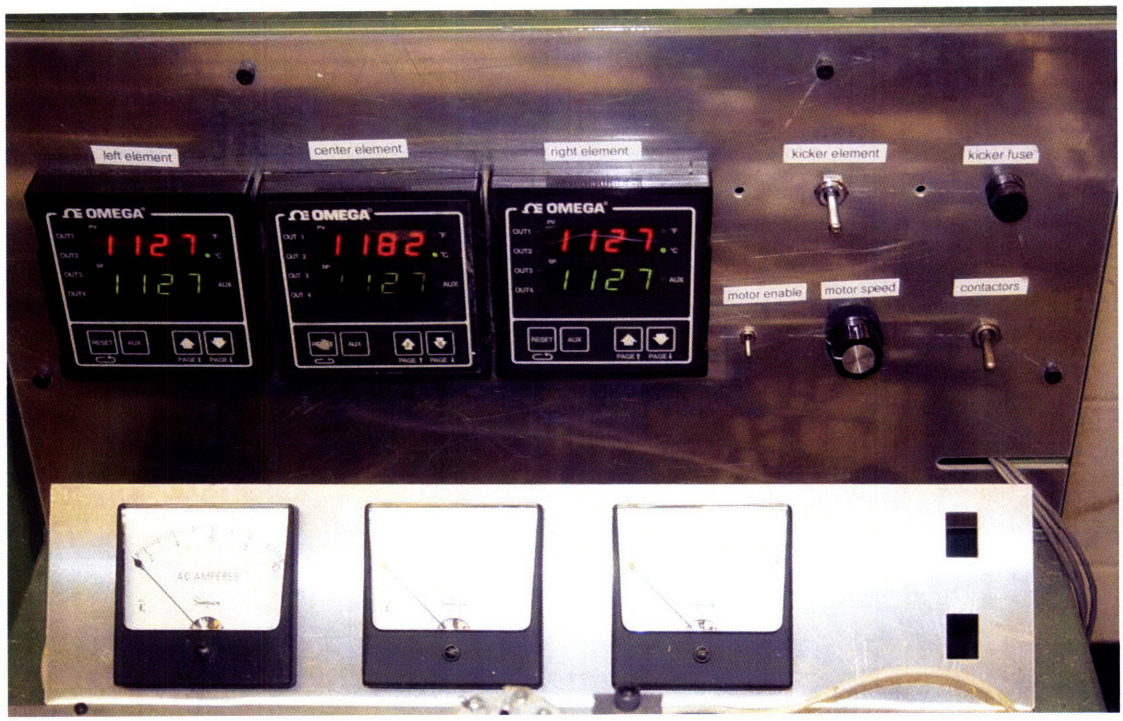

Figure 12: Control panel of the completed furnace. An Omega CN3251 (top) and an analog AC ammeter (bottom) are assigned to each loop. Analog meters provide continuous readout of current, without aliasing – this becomes important during fast transients.

4.3 Insulation

Furnace Energy Payback

In order for the recrystallization process to make economic sense, the energy expended in performing it must be delivered by the finished solar cell on a short timescale relative to the cell's 25-year lifespan. An energy payback period on the order of a few months is a reasonable goal for this process step. Accordingly, given conservative baseline assumptions of installed cell performance, it is possible to establish maximum power consumption standards for the recrystallization furnace. Assuming the following PV system performance:

12% installed module efficiency

1000 W/m^2 incident at peak sunlight conditions

1000 Wh/yr/W_p installed

Cell length = 12cm

Furnace traverse time = 0.1 h (6cm/min feedrate)

Each centimeter width of the installed solar cells will yield

0.0012 m^2 × 1000W/m^2 × 0.12 (Efficiency) = 0.144 W_p.

Accordingly, each cm width of installed solar cells will yield

0.144 Wp × 1000 Wh/yr/Wp = 144 Wh/yr = 12 Wh/month.

For one-year energy payback, this performance corresponds to a furnace power of

144 Wh / 0.1 h = 1440 W / cm width;

for one-month payback, furnace power must be

12Wh / 0.1h = 120 W / cm width.

Given more generous system performance:

14% installed module efficiency

1000 W/m^2 incident at peak sunlight conditions

1500 Wh/yr/W_p installed

Cell length = 12cm

Furnace traverse time = 0.1 h (6cm/min feedrate)

Each cm width of solar cells will yield 0.168 Wp, corresponding to an annual yield of 252 Wh/yr = 21 Wh/mo. This would require 2520 W / cm width for one-year payback, or 210 W / cm width for one-month payback.

Thus, minimizing heat losses and insulating effectively are essential to rendering the recrystallization process worthwhile. These efficiency drivers necessitated experiments in radiant shielding and insulation materials, which are outlined below.

Insulation Material Choices

The choice of insulation materials in a recrystallization furnace is constrained by contamination risk mitigation (impurity diffusion takes place quickly at recrystallization temperatures), effectiveness as a radiation shield, and effectiveness as a conductive thermal insulator.

Mitigating the risk of contamination requires that compounds containing either transition metals or dopants are excluded. Graphite is ruled out for its lack of sufficient temperature capability in air. Consequently, for the temperature range in question, the two most suitable compounds are Si_3N_4 and SiC.

However, SiC and Si_3N_4 are typically only available in dense (i.e., not insulating) forms. Furthermore, most commercially-available SiC and Si_3N_4 include transition metal impurities. In order to make functional insulating material out of these compounds, it was necessary to form porous blocks by sintering high-purity powders of stock material together into solids. As KION, Inc.'s polysilazane SiC precursor (used as a binder) and a multitude of different SiC powders was readily available, SiC was used as a baseline insulation material.

Sacrificial materials were introduced into the mixture of KION and SiC powder in order to decrease the density of the sintered solid and thus increase thermal resistance. Initial tests used sugar, crushed walnut shells, and graphite powder as a sacrificial porogen; the mass of sacrificial material was subsequently removed by oxidation. The formulation of insulation could be fine-tuned depending on the porogen used and the proportion of porogen mass to the overall mass of the mixture.

It was desired to characterize both the radiation shielding and thermal conduction insulating properties of the insulation materials under development. By evaluating the relative impact of porosity and insulation thickness, the formulation of insulation could be more finely tuned. Accordingly, a test furnace was constructed out of silica-alumina refractory ceramic fiber, using IsquaredR's Starbar heater elements. The furnace had a 51mm x 51mm square aperture for mounting insulation samples. Samples were compared based on the furnace power required to maintain a temperature of 1455°C, as measured by a thermocouple situated between the furnace's heater elements.

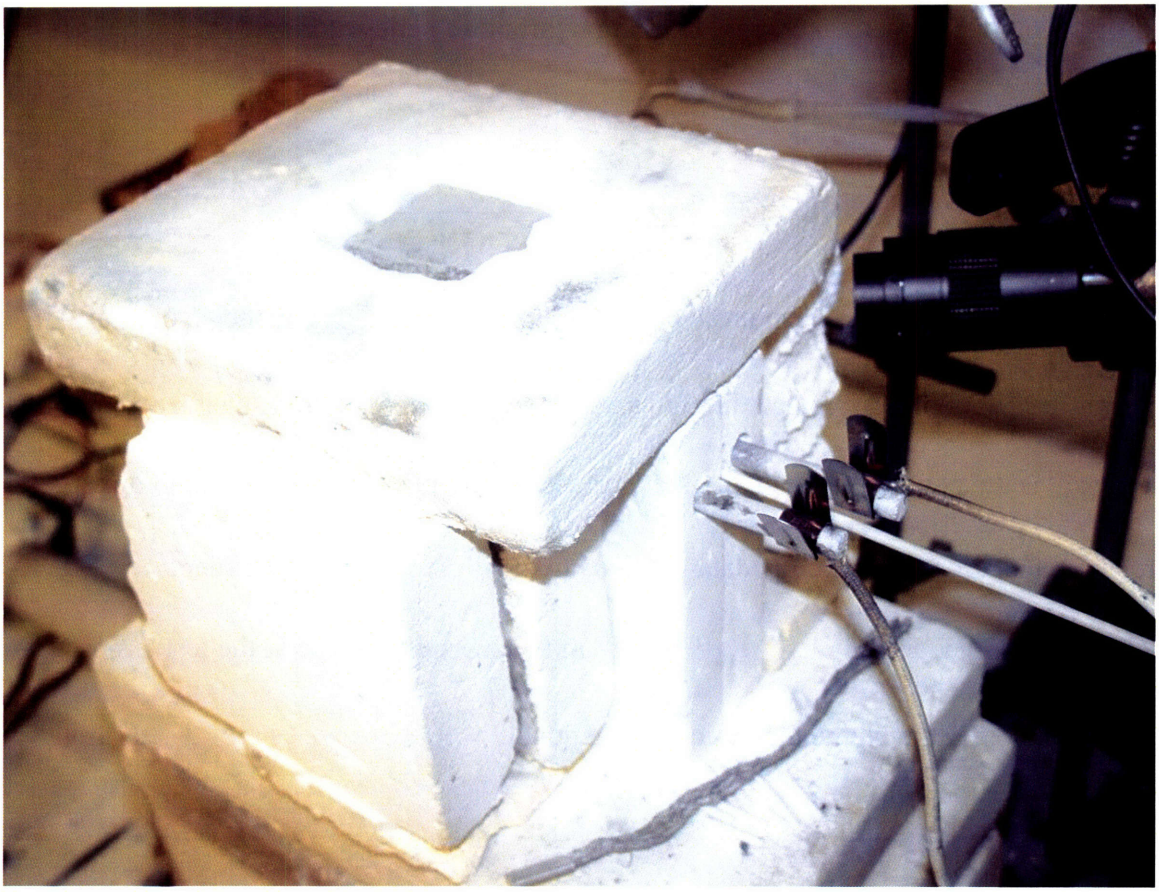

Figure 13: Aperture furnace used for insulation testing. The material under test sits in the port in the furnace lid. The heater elements are IsquaredR's Starbar silicon carbide heaters. A type R thermocouple is used for temperature feedback.

Two experimental controls were used. The first control was intended to evaluate the coupled radiative and conductive insulating properties of readily available refractory

insulations. The second control was intended primarily to evaluate the merits of purely radiative shielding.

For the first control, a fitted block of off-the-shelf silica-alumina refractory ceramic fiber insulation (the same material as the rest of the furnace) was placed in the aperture. For the second control, the refractory ceramic fiber insulation was removed and a stack of monolithic dense Poco Hexoloy SiC tiles were stacked sequentially in the aperture of the furnace. The control tiles were 1.25mm in thickness and 51mm square, spaced apart by 3mm thick alumina spacers.

The rate of thermal energy dissipation through a series of radiation shields is given by

$$\dot{Q}_{total} = \dot{Q}_{elsewhere} + \frac{\dot{Q}_{open\_aperture}}{n+1}$$

where is $\dot{Q}_{total}$ is the total amount of power dissipated, $\dot{Q}_{open\_aperture}$ is the amount of power that would be dissipated through an aperture with no radiation shields in place, and $\dot{Q}_{elsewhere}$ is the power loss from all other surfaces of the furnace. According to this relation and a regression of the data gathered in the SiC tile stack experiment, $\dot{Q}_{elsewhere}$ = 470W and $\dot{Q}_{open\_aperture}$ = 1000W in the aperture furnace at the test temperature. The data closely match the theoretical prediction.

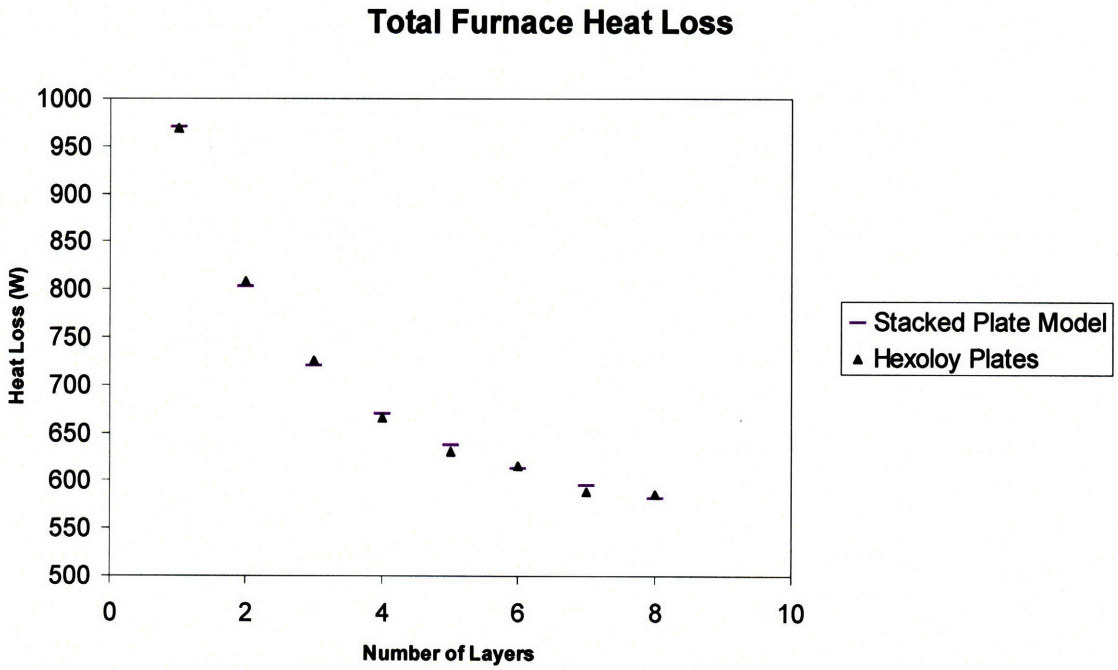

Figure 14: Heat loss through a stack of radiation shields. The model (in red) matches the measured values (blue) closely, allowing precise calculation of the $Q_{elsewhere}$ term.

From this point, it is possible to compare the relative performance of the different formulations of SiC insulation that were developed.

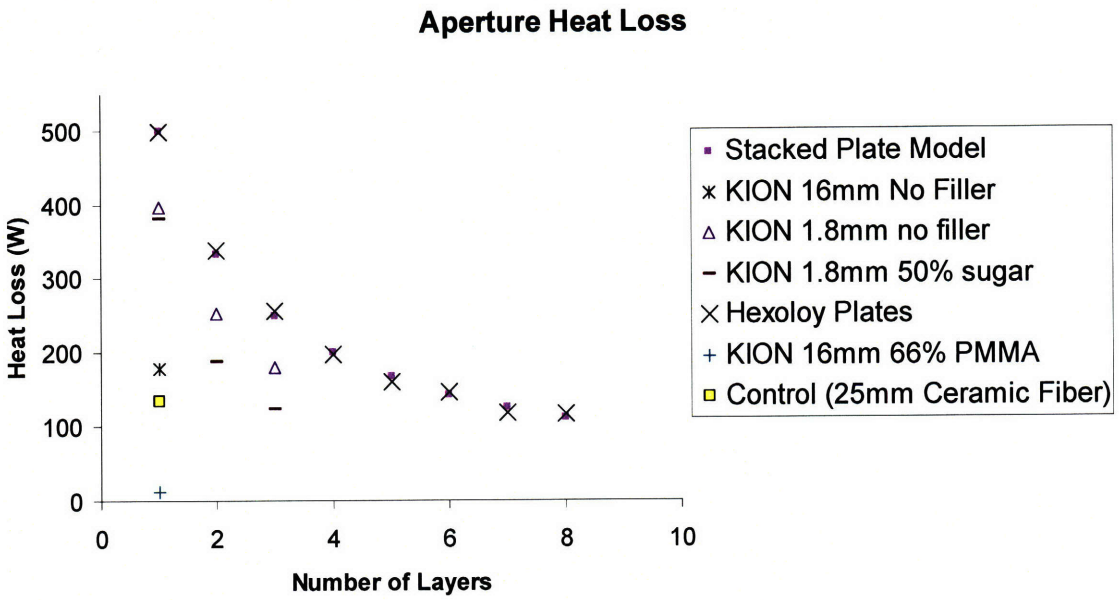

Figure 15: Performance of different insulation formulations tested, superimposed on the plate model and theoretical prediction plot from above. KION 16mm 66% PMMA represents the best formulation of insulation developed; it was used in the final furnace design. It outperformed off-the-shelf refractory ceramic fiber insulation.

Because this experiment sought to test the influences of porosity and thickness (and therefore the radiative vs. conductive thermal resistance of insulation materials), samples with and without filler, as well as thick and thin samples were tested. From Figure 15 above, several samples happen to fall around 180W for their aperture heat loss – meaning that their thermal resistance is roughly equivalent. Thus, a thick piece (16mm) sample of unfilled KION insulation and SiC powder; a stack of 2 filled, thin (50% dense, 1.8mm) KION and SiC plates; a stack of 3 unfilled, thin KION (100% dense, 1.8mm) and SiC plates; and a stack of 4 to 5 fully dense Hexoloy (1.3mm) plates offer the same thermal resistance.

Thus, one can conclude that the number of plates, plate thickness, and plate porosity all play a non-negligible role in the effectiveness of insulation. From the data, it is apparent that insulation of the same formulation has equal thermal resistance when in a thick single layer (16mm) or in 3 distinct layers (each 1.8mm). Thus, strictly from the

standpoint of simplicity of design, it is preferable to cast a single thick insulation wall rather than multiple thin walls arranged in series. The thermal resistance benefit of adding radiative gaps in this material was not sufficient to make it an interesting design possibility over simply thick walls.

In reality, porosity confers similar benefits to adding radiative gaps in the material. With sufficiently porous material, conductive paths are short (on the order of constituent particle size) and heat must make radiative jumps to pass on to successive solid particles in the material. However, given the data above, the benefits of porosity are not realized with thin walls. For instance, heat loss is only improved by 5% by increasing porosity to 50% in 1.8mm samples. The benefit is much more pronounced with 16mm samples: heat loss is cut effectively to zero (as measured using the present technique) by increasing the porosity of a sample from 0 to 66%.

Finally, insulation thickness is significant, but only up to a point. Returns on thickness increase begin to diminish as the system pushes the limit of losing as much heat through increased exterior surface area as it is gaining through conductive thermal resistance. Ultimately, insulation thickness was chosen on the basis of what would fit within a reasonable physical footprint in the space available. 25mm insulation thickness was employed in the furnace.

Similarly, while it was informed by experimental results, insulation formulation was ultimately chosen on the basis of physical implementation. As porogens, sugar, graphite, and walnut shells were ineffective (yielding weak and flaky finished parts), likely because residue from the oxidized powders impeded full sintering of neighboring SiC granules. For this reason, acrylic beads (PMMA) were used as a porogen. Above

460°C, acrylic oxidizes completely, yielding only CO_2 and water as combustion products. Ultimately, a 2:1 mixture (by mass) of 30μm acrylic beads, 2μm SiC, powder, and KION yielded insulation blocks of sufficient mechanical robustness to hold together over the life of the furnace.

Resulting Furnace Performance

With the insulation material employed in the final furnace design, steady-state power draw of the entire furnace in recrystallizing a 2.8 cm wide sample was approximately 1500W, or 535W/cm width at a pull rate of 0.6cm/min (one tenth the feedrate used for the above analysis). Increasing feedrate and substrate/backing plate length will increase power draw per unit width, but potentially only marginally. Most heat loss is out the doors of the furnace, the rate of which will not be affected by higher sample pull rates. Current performance is well within the range of one-year energy payback. As such, prospects are promising for realistic energy paybacks when scaled up.

4.4 Motion System

In order to move the sample through the furnace, a linear motion slide was installed through its length. Two longitudinally sliding bars oriented parallel to the direction of sample travel provided a platform on which to set samples. The slider bars were made of Saint Gobain Hexoloy SE, as were the bearing surfaces on which the slider bars slid. Due to sliding surfaces and inevitable particle generation, purity of these components is particularly important.

The motion system was designed with sufficient travel to allow a sample to pass completely through the length of the furnace in a single motion, in order to ensure repeatability from run to run.

Motion was actuated with a leadscrew drive, using a DC tachomotor from Faulhaber. The motor's integral tachometer was wired into a Faulhaber MCDC3003 motor controller, for closed-loop speed control. The system was capable of forward and backward motion, from 0 to 80mm/min in both directions. All moving components were mounted to a linear slide, mounted on linear bearings which bore on dual Thompson shaft guide rails.

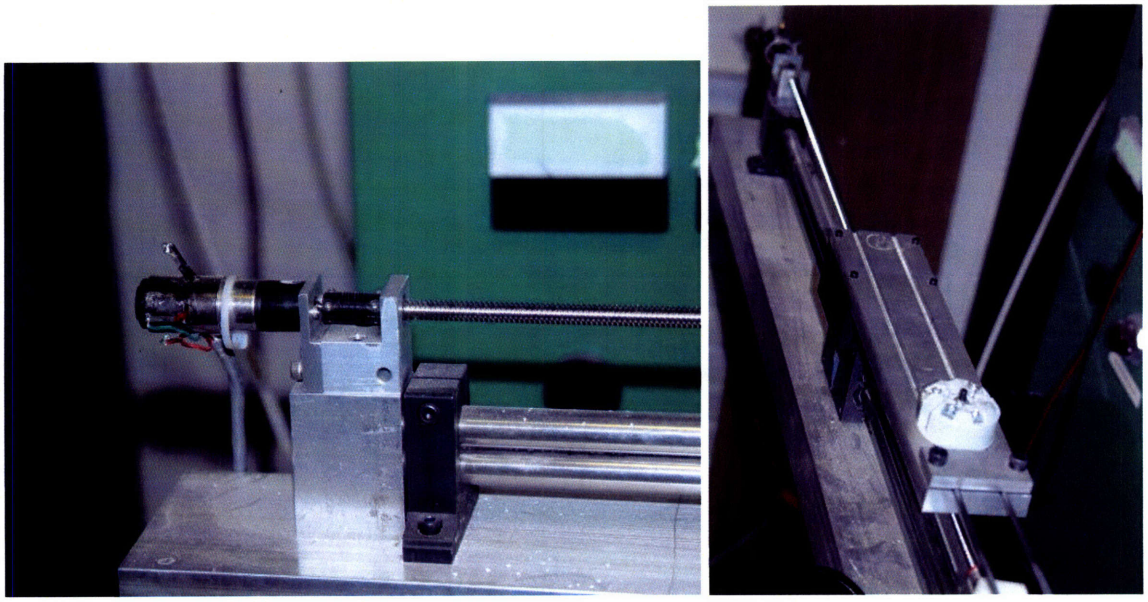

(a) (b)

Figure 16: Tachomotor and lead screw (a) and entire drive system assembly (b). The Hexoloy slider bars are visible in the lower right corner of the picture on the right.

4.5 Final Furnace Implementation

The furnace was implemented in accordance with the development projects outlined above. Below is an image of the final design used to recrystallize all the samples used in recrystallization experiments.

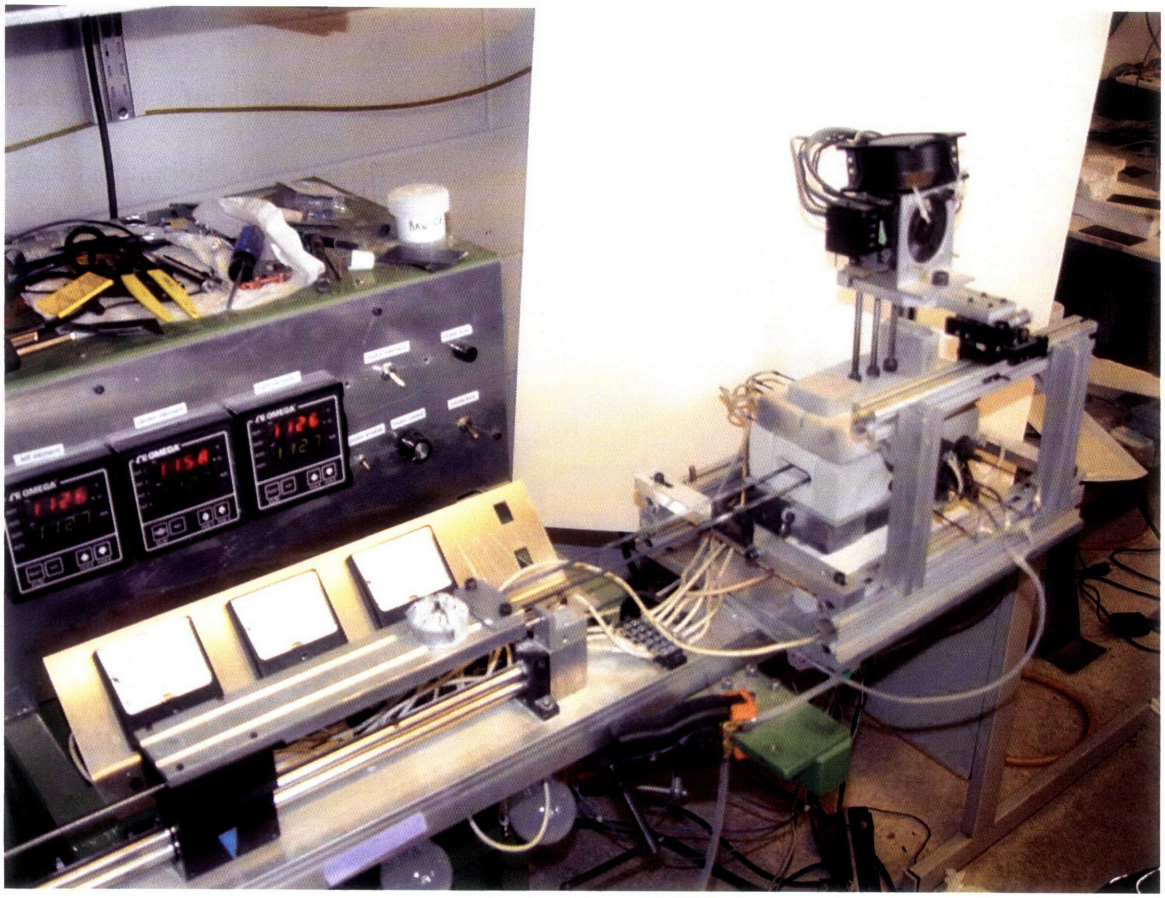

Figure 17: Final furnace implementation.

Chapter 5: Thermal Profile Experimental Procedures

Following the set of initial proof-of-concept recrystallizations outlined in Chapter 3, a set of samples was processed in order to evaluate the effects of recrystallization on grain structure and dislocation density, particularly given differing temperature profiles.

5.1 Samples

Stock

Standard cast multicrystalline silicon was used as a feedstock for recrystallization experiments. Starting wafer thickness was 300μm, with the overall wafer dimensions 28mm in width by 46mm in length. Typical starting grains were between 1mm and 10mm as measured along their largest dimension.

The rapid solidification casting technique sought as a precursor step to the recrystallization of this experiment will likely yield small-grained wafers (with grain size on the order of microns). As such, for proof-of-concept experiments, it would be ideal to use microcrystalline wafers as feedstock.

However, given the lack of a supplier for microcrystalline wafers, this study employed standard multicrystalline silicon. The reasoning was that, given complete melting, recrystallization behavior should be similar regardless of the grain structure of the starting wafer. Thus, assuming no memory effects from the oxide capsule, cast multicrystalline silicon can function as an effective model for cast microcrystalline silicon.

Late experiments sought to transition to Motech's standard cast multicrystalline silicon feedstock (180μm thick), but its thinness presented handling difficulties (which only became more pronounced after recrystallization). As a result, the thicker 300μm stock was maintained for baseline experiments.

Preparation

For thermal profile experiments, wafers were sawn down to a 28mm x 46mm sample size on a diamond die saw with UV-releasable dicing tape backing. Following release from the dicing tape, samples were sonicated for 5 minutes in N-methyl pyrrolidone to strip them of residual particulate material and organic contaminants. Subsequently, samples were etched for 60 seconds in a 9:0:1 isotropic etch to remove preexisting saw damage.

Next, the containment capsule for recrystallization was grown. While several options and many variations thereof exist for types of capsule, this study simply chose a functional baseline and used it across all experiments. The baseline capsule was a dry thermal oxide, grown in a quartz tube furnace over 20 hours at 1100°C in air. The resulting oxide capsule was approximately 600nm thick.

At all steps of the process, samples were photographed (front and back) to allow comparison from step to step.

5.2 Backing Plates

Stock

Saint Gobain's Hexoloy SE extruded silicon carbide plate was used as a substrate backing, for both the top and bottom of the sample. Thickness of the backing plates was

1.25mm. Plates were sawn down to the largest possible size that could fit through the recrystallization furnace: a length of 51mm and a width of 30.5mm. Samples were cut from plate stock using a wet-blade diamond saw.

Release Layer

Development of the release layer system is detailed in Chapter 3. The release layer used in all thermal profile experiments was applied with the thin-layer adhesion and sintering method, outlined in 3.2. This technique provided reliable, repeatable release and a flat, highly-dense surface layer.

5.3 Furnace Setup

Alignment

Before each run, several components in the furnace were checked for alignment. These components are described in greater detail in Chapter 4. First, the housing of the furnace was adjusted to ensure that sample travel was perpendicular to the heater elements and directly centered in the left-right direction of the furnace. Furthermore, the entire furnace housing was shimmed up or down in order to maintain the desired top-bottom spacing of the sample between the heater elements. Finally, the pyrometers were checked for alignment directly above the heater elements. With the main body of the pyrometer detached, a laser pointer was used to direct a beam through the collimating optics and onto the target elements. The pyrometers were mounted on a 2-axis micrometer stage (which allowed adjustment in X and Y) in order to ensure that the pyrometer's sampling spot was directly in the center of each element. The laser pointer provided simple visual feedback for adjusting the pyrometers' micrometer stage.

In the future, more attention should be paid to the dimensional coupling between the heater elements and the slider bars. Typically, alignment was optimized for uniform spacing of the sample from the heater element without regard to levelness of the slider bars. As a result, the slider bars held the sample at a slight tilt during melting, which potentially could have impacted thickness uniformity of recrystallized samples.

Temperature Setpoints

Before doing any recrystallizations in the furnace, a profiling thermocouple was used to determine the sensitivity of temperature profile to element temperature setpoints. To generate the profiles, a Type R thermocouple was clamped to the sample slider rails and passed through the furnace at 3.6cm/min. The thermocouple was wired to an Omega TX13 thermocouple transmitter, which output a corresponding analog signal to a Tektronix datalogging oscilloscope. The 3.6 cm/min feedrate corresponds to the slowest feedrate that still fits within the widest time window of the oscilloscope's datalogger.

While the background temperature experienced by the thermocouple was predictably lower than the heater element setpoint, attempts at locally fine-tuning the profile yielded minimal change. To quantify the effects, temperature profiles (left to right) were set at 1575°C, 1127°C (the controller's minimum resolvable temperature), and 1575°C.

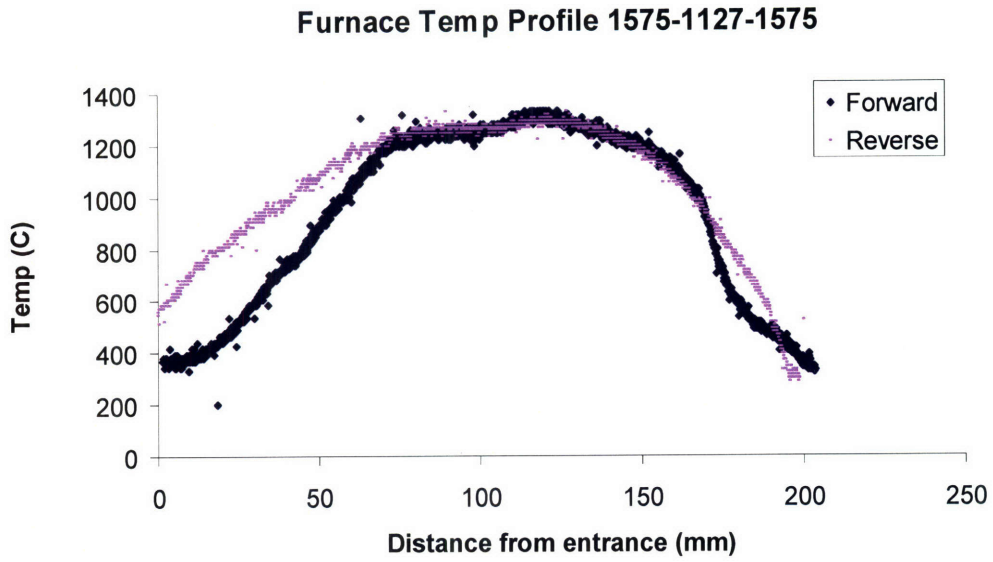

Figure 1: Temperature profile taken to examine the effects of varied heater element temperature setpoints. The bold points are from a run in which the thermocouple was passed through the furnace forward; the narrow points are from a run in which the thermocouple was pulled through the furnace backward. Since reverse pulls send heat from the thermocouple sheath into the bead, the effect is a rounded profile that does not reflect true furnace temperature profile. The best fidelity is achieved during a forward feed. Accordingly, all subsequent temperature profiles were taken feeding forward.

As is apparent above, even though the center element (located 100mm from the entrance of the furnace) was turned completely off, the resultant dip in temperature profile was small compared even to offset error between pyrometers. As such, it was decided to run recrystallization experiments with a constant temperature setpoint from element to element, and to effect changes in temperature profile by moving blocks of insulation around.

In all recrystallization experiments, heater elements were set to a temperature of 1625°C, as read by the control pyrometer. The exception was the 3 recrystallization runs under the gentle profile, in which the leading element was run at 1655°C, as lower element temperatures failed to actually melt the sample.

Insulation Geometry

Depending on the temperature profile sought, the exit door of the furnace was either relocated or removed. The exit door was left in place to achieve a standard profile, was moved closer to the center of the furnace for a steeper exit profile, and removed altogether for a gentle exit profile.

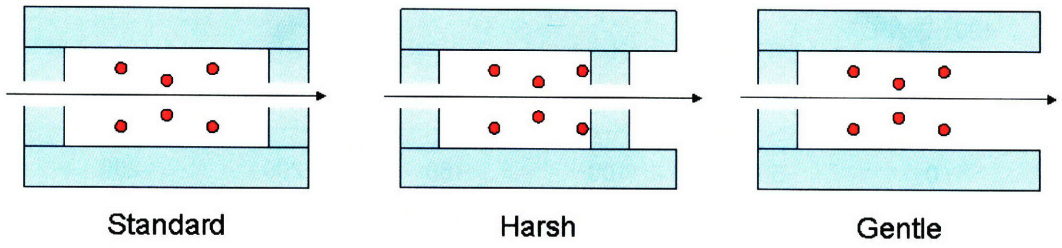

Figure 2: Cross-sectional diagrams of the 3 insulation arrangements employed to achieve different thermal profiles. Sample travel path is indicated by the arrow in the sketch.

In order to test the resulting temperature profile, a profiling thermocouple was affixed to the slider bars and passed through the furnace, attached to a datalogger. The resulting profiles are indeed macroscopically different.

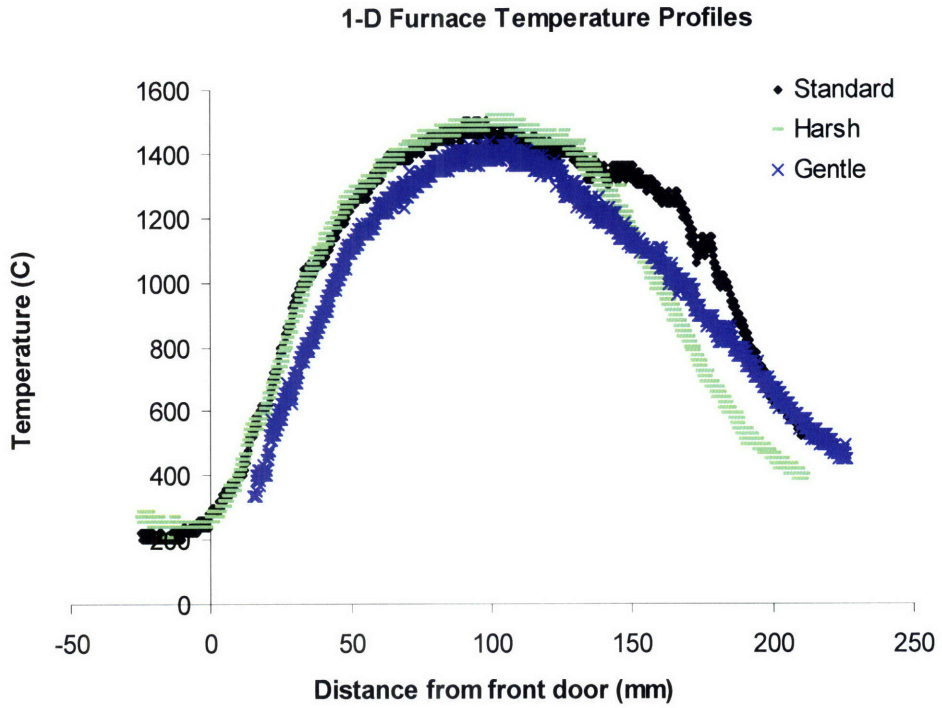

Figure 3: Longitudinal furnace temperature profiles, as determined by profiling with a moving thermocouple.

5.4 Thermal Process

Preheating

Given the high resistivity of the Hexoloy heater elements at room temperature, a preheat cycle was performed in order to start the main elements (see Chapter 5). Once conducting on their own, the furnace (due to the conservative control parameters) was allowed approximately 20 minutes to stabilize at its temperature setpoint.

Recrystallization

Each sample was assembled into a stack with backing plates on top and bottom. Once assembled, it was set down on the slider rails directly outside the entrance plane of the furnace. The long direction of the sample was oriented parallel to the direction of

travel through the furnace. From there, the sample was driven through the furnace. As the effects of pull speed are beyond the scope of this experiment, all samples were moved at 6mm/min, an existing standard pull speed from earlier proof-of concept work.

Once completely clear of the exit plane of the furnace, samples were removed from the slider rails and allowed to cool to room temperature.

5.5 Post Processing

Oxide Etch

Given the thick layer of oxide grown on the sample, it was necessary to perform an oxide etch to allow subsequent processing steps. Samples were placed in a 10:1 solution of deionized water and 49% hydrofluoric acid for 30 minutes. The endpoint of etching was verified by observing beading of water over the whole area of each sample.

Texture Etch

To reveal the grain structure of the samples, each was placed in an anisotropic etch solution of hot concentrated sodium hydroxide. Sodium hydroxide etches preferentially etch the <100> direction of silicon crystals, which yields readily visible patterns in the surface of the sample: grains are visually differentiable.

The etch solution was held at 80°C in a water bath on top of a hotplate. Deionized water was heated to 80°C in the water bath, and then sodium hydroxide flakes were added until saturation (i.e., further flakes failed to dissolve). Samples were etched for 3 minutes in this solution.

Polishing

In order to prepare for the subsequent dislocation etch, samples were ground and polished. This process ensured that etch pits contrasted uniformly with the same dark background from sample to sample.

Samples were mounted to thick (2.5cm) aluminum blocks, using black sealing wax on a hotplate. The blocks provided a sturdy support for handling during polishing, and ensured the samples were not loaded in bending during any step.

Figure 4: Samples mounted to aluminum blocks with wax, before grinding and polishing.

Grinding and polishing was carried out on two pieces of hardware. First, a 5-inch diamond-impregnated aluminum grinding wheel was used to rough-cut the top surface of the sample into a flat plane. From that point, a Buehler Ecomet IV polishing wheel was used with silicon carbide abrasive paper to wet-polish the top surface of each sample to a 1200 grit finish. At this level of surface finish on silicon, residual scratches from polishing are on the order of 0.2-1μm. Given that a planarizing etch precedes the dislocation etch, and that the size of the dislocation etch pits is 3-5μm, a 1200 grit finish was sufficient for resolving etch pits.

After polishing, samples were removed from their aluminum mount blocks by reheating the aluminum blocks on a hotplate until the mounting wax melted. Following release, samples were stripped of residual wax with a toluene soak, followed by an acetone-ethanol-deionized water rinse.

Dislocation Etch

In order to quantify dislocation density, a Sopori etch was used to mark surface dislocations with etch pits. A 3-step process was used to optimize etch pit clarity. After washing, each sample was placed in 9:0:1 planarizing etch for 60 sec, then transferred directly to the Sopori etch solution for 30 seconds after the first formation of bubbles, and then quenched momentarily in 9:0:1 etch to prevent re-plating of silicon onto the sample. After a final quench in deionized water, samples were ready for imaging.

The extent of dislocation etching is sensitive to several parameters: etch bath temperature, etch composition, and time spent in the etchant. To ensure repeatability between samples, large batches of 9:0:1 and Sopori etch solution were mixed. Each

sample was etched in freshly poured beakers of etchant; each dose of etchant was used for only one sample.

Furthermore, time spent in the etchant was optimized for the bright vs. dark method of quantification described below. In areas of high dislocation density, closely-spaced etch pits (bright spots) can overlap with one another into a uniform trench, giving the appearance of darker, non-dislocated material. To avoid this misleading effect, residence time in the dislocation etch was optimized.

Unrecrystallized samples of cast multicrystalline silicon were damage etched, cleaned, and Sopori etched in the manner described above, for 15, 20, 25, and 30 seconds, respectively. Areas of high dislocation density from each sample were imaged in darkfield mode on an optical microscope.

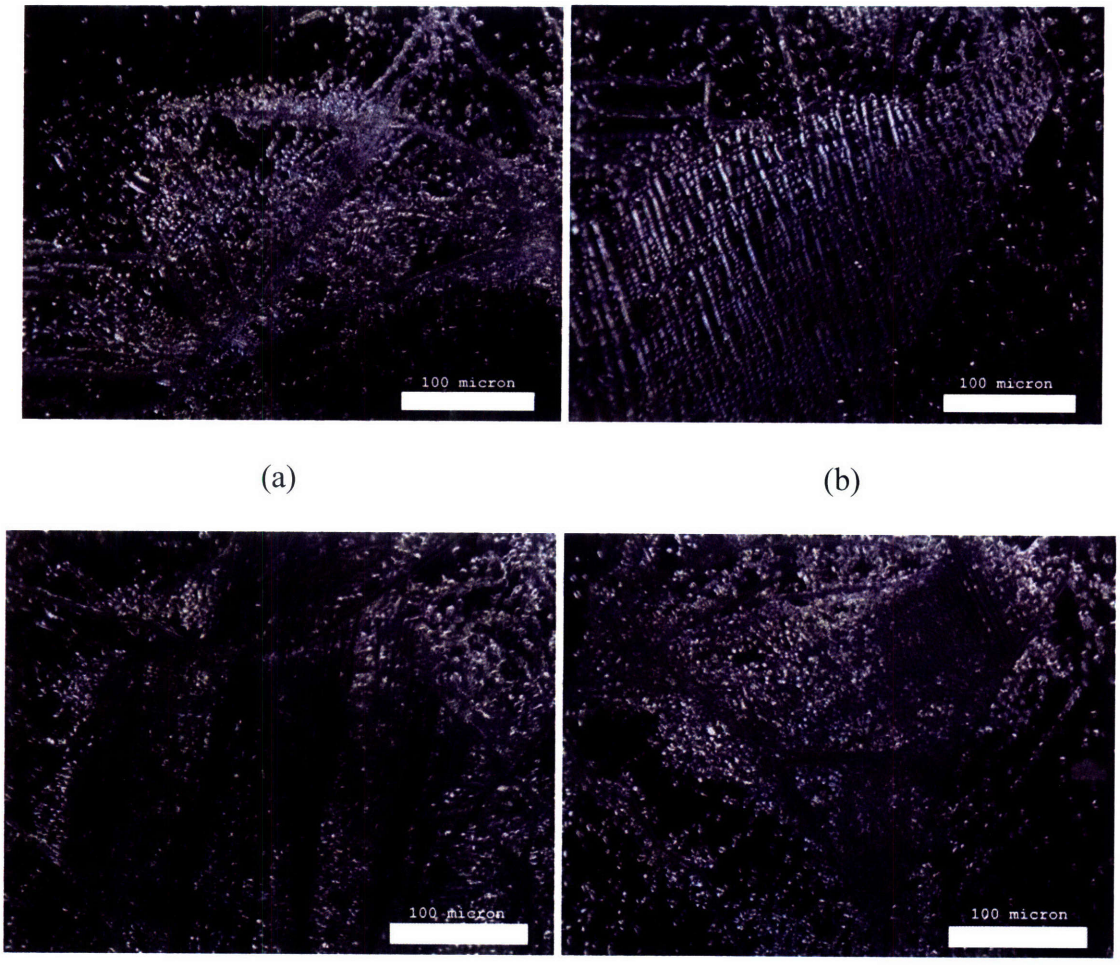

Figure 5: Optical micrographs of 15 sec (a), 20 sec (b), 25 sec (c), and 30 sec (d) etches, respectively. Beyond 20 sec, the etch pits begin to overlap, causing dark striations. As such, 20 sec was the decided upon time for Sopori etching.

As is apparent in the micrographs, neighboring etch pits in areas of high dislocation density begin to overlap with one another at etch residence times above 20 sec. This apparent cancellation of etch pits manifests itself as areas of dark striations emerging in a smooth gradient from light-colored areas. From this experiment, it becomes clear that 20 sec is the maximum allowable etch time. Since it was also desired to maximize the appearance of dislocations, the maximum allowable time was used; all samples were dislocation etched for 20 sec.

Imaging and Analysis

After dislocation etching all samples, each was placed face-down on a Hewlett Packard Scanjet 5470c flatbed scanner and scanned into a .tif file at 2400 dpi resolution.

Figure 6: Example of the as-scanned surface, post dislocation etching. Dislocation pits appear as light spots on the surface.

In the resulting images, dislocation pits appear as light colored spots. Thus, in order to quantify the extent of dislocation across the entire sample, the pixels of each image (after bring cropped and converted to grayscale in Adobe Photoshop) were plotted in a histogram. Ranging from dark (0) to light (255), the values at each pixel were binned and their frequency of occurrence was plotted. By comparing the distributions of dark and light pixels, a quantitative map of dislocation density could be charted from sample to sample.

Chapter 6: Results

Overall, recrystallized samples in the standard and gentle temperature regimes showed decreases in dislocation density and decreases in minority carrier lifetime relative to controls. Samples recrystallized in the harsh temperature regime increased in dislocation density while also decreasing in minority carrier lifetime. However, the decrease in lifetime for harsh samples was less pronounced than the decrease in gentle and standard sample lifetimes.

In all temperature profiles, resulting grains were several millimeters in width, and tended to be long (some nearly the full 46mm of sample length) in the pull direction of the furnace. Of the three temperature profiles tested, samples recrystallized under the standard temperature profile yielded the largest grains and lowest dislocation density.

Photoconductance of recrystallized samples was significantly lower than that of control samples; the difference in photoconductance in part explains the decreases in measured minority carrier lifetime.

6.1 Grain Structure

Starting Material

The cast multicrystalline silicon feedstock began with grains of several mm to cm in principal dimension.

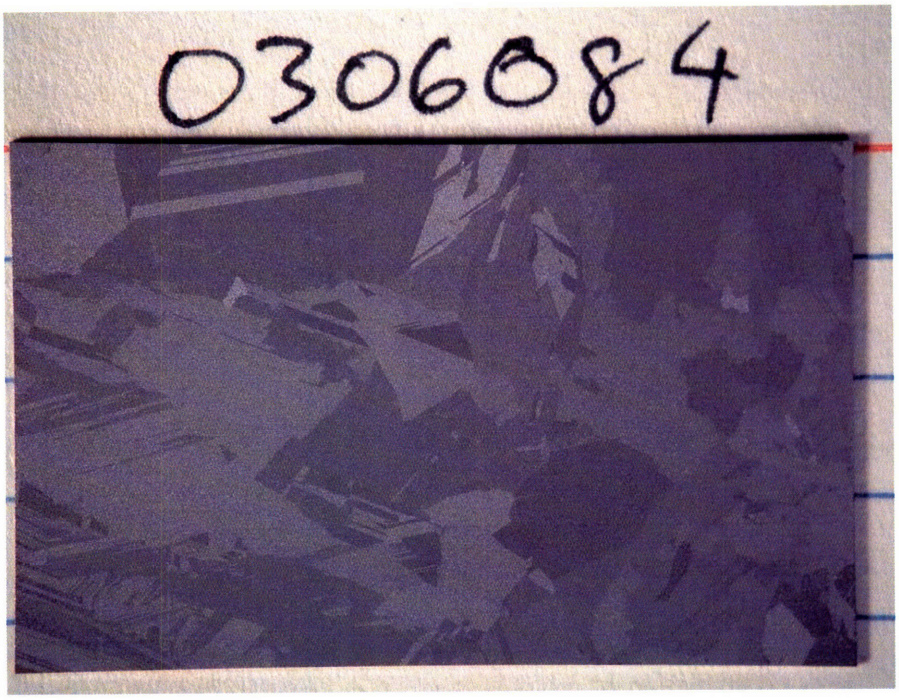

Figure 1: Typical starting wafer after cleaning. Overall sample size is 28x46mm.

Grain Structure After Recrystallization

In all recrystallized samples, ending grain size was on the same order as starting grain size. Size ranges from millimeters to centimeters were typical. Grain shape tended to be elongated, in the direction of travel through the furnace. Many grains extended to nearly the entire length of the sample, with grain boundaries tending to run parallel and at shallow angles (approximately 30°) to the pull direction.

Many grain boundaries tended to be perpendicular to the plane of the sample, as opposed to parallel. (Grain boundaries parallel to the plane of the wafer are particularly damaging to PV device performance.) Top-bottom views of a given sample show many similarly-shaped grains present, which is an indicator that most grain boundaries are not midplane. However, large samples typically showed less front-to-back uniformity than small samples taken from the same stock material; these are outlined in 3.4. The cause for this difference is in need of further investigation.

Figure 2: Typical grain structure after recrystallization and texture etching. Above is a top view of a sample, and below is a bottom view of the same sample, rotated about its long axis. Grains are millimeters to centimeters in length, with elongation down the pull direction of the sample and at a shallow angle (30° to the pull direction). The differences in surface finish obscure some grains, but most features are still visible.

Dependence on Temperature Profile

Samples recrystallized in the harsh and gentle temperature profiles tended to have narrower grains than samples recrystallized in the standard temperature profile. Lengths of grains were comparable.

The following pages contain photographs of all samples recrystallized in each temperature profile. Samples have been texture etched to reveal the underlying grain structure.

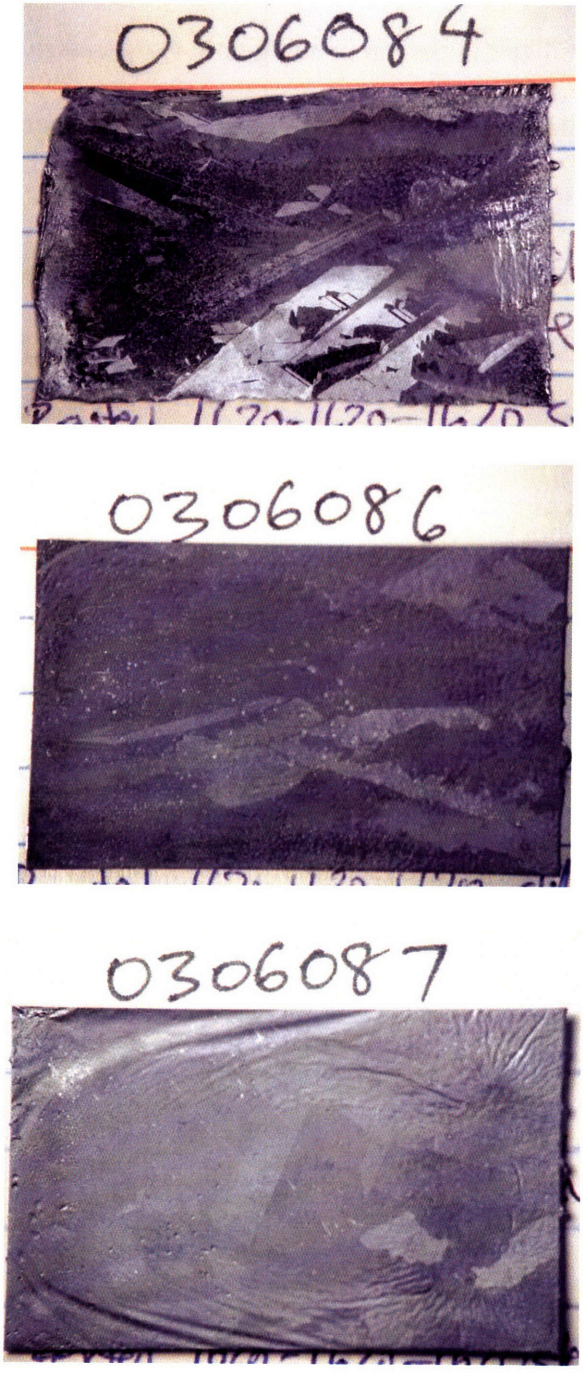

Figure 3: Top sides of samples recrystallized under the standard temperature profile, after texture etching. Surface finish varies significantly from sample to sample, although grain structure is still visible. These differences in surface finish stem from an incomplete oxide etch prior to texture etching. As these three samples were the first set recrystallized, some processing errors were committed.

Figure 4: Samples recrystallized under the harsh profile, after texture etching. The furnace failed when sample 0407083 was only partway through, so only the right half is recrystallized. The successfully recrystallized half was still used as a data point.

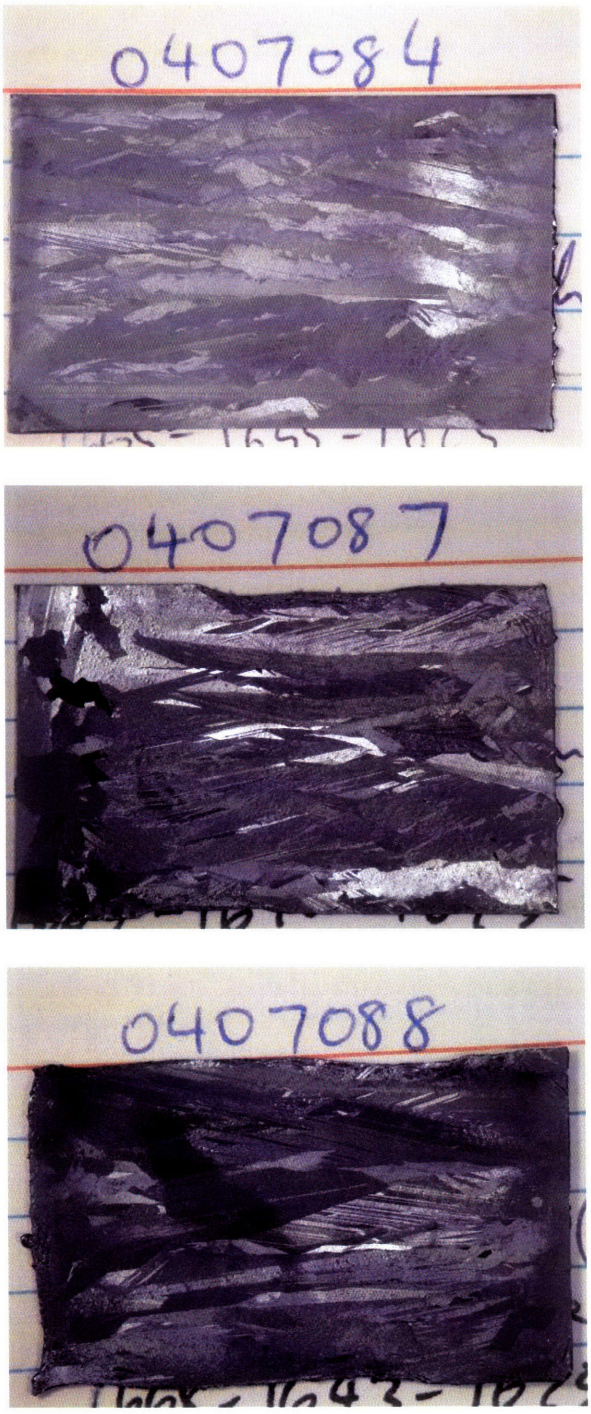

Figure 5: Samples recrystallized under the gentle profile, after texture etching. 0407087 did not fully melt; the left edge was digitally cropped out for the dislocation density study.

In examining the temperature peaks of the furnace profiles of 5.3, the profiles the sample experiences in gentle and harsh runs are actually the most sharply concave down – particularly when examining the segment above 1000°C, at which the sample is just beginning to refreeze. This increased concavity $\frac{\partial^2 T}{\partial x^2}$ of the temperature profile leads to increased thermal stress as explained in 2.2. In turn, high thermal stress leads to dislocation formation and propagation in the growing crystal. Dislocations at the freeze front are capable of nucleating new grains, as explained in 2.3. This behavior is one possible explanation for the smaller, narrower grains visible in the gentle and harsh samples.

6.2 Dislocation Density

Starting Material

Material entering the process had peak dislocation density (at grain boundaries) of 8E6 cm^{-2} resolvable dislocations, with a typical bulk dislocation density of 8E5 cm^{-2}. Dislocation densities were measured by imaging on an optical microscope, finding representative areas of maximum and minimum dislocation density, and manually counting etch pits. These densities were in turn mapped to peak and mode brightness values in area scans of wafers, as described in the following section.

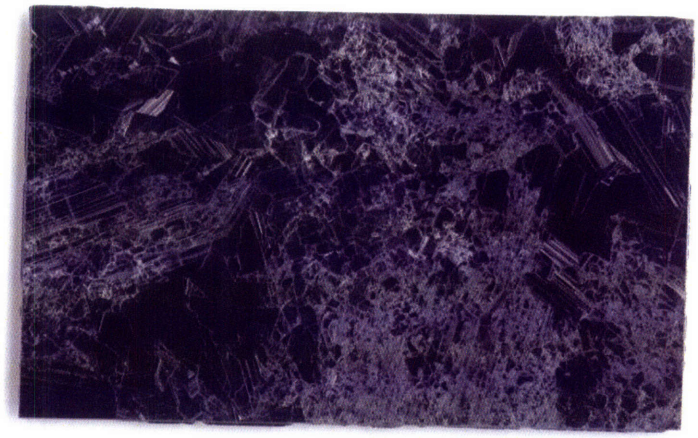

Figure 6: Flatbed scan of dislocation etch pits on an unrecrystallized sample. The sample was subject to the same damage etch that all recrystallized samples were subject to at the start of processing.

Dislocation Density After Recrystallization

Recrystallized samples showed the following bulk and peak dislocation densities:

Profile	**Bulk Dislocation density**	**Peak Dislocation density**
Standard	4E5 cm^{-2}	2E6 cm^{-2}
Harsh	8E5 cm^{-2}	1E7 cm^{-2}
Gentle	6E5 cm^{-2}	2E6 cm^{-2}
Control	8E5 cm^{-2}	8E6 cm^{-2}

Dislocation densities were determined by mapping scanned shades of dark and light to micrographs in which etch pits were manually counted. Once scanned, a histogram of pixel brightness (0-255) for each sample was created, with brightness values binned into 128 bins. From these histograms, dislocation density distribution was established: peak

dislocation density was determined by the brightest populated bin of the histogram, and bulk dislocation density was determined by the bin with the highest population of pixels.

Peak dislocation density tends to correspond to grain boundaries, where dislocations accumulate. Thus, there is some dependence of measured dislocation density on grain boundary density. This dependence is apparent in the histogram at the end of this section.

The following pages contain the scanned images used for dislocation density measurement.

Figure 7: Scans of polished and Sopori-etched samples, recrystallized under the standard profile.

Figure 8: Scans of polished and Sopori-etched samples, recrystallized under the harsh profile. Sample 0407083 was only partially melted due to a heater failure; the right half of the sample was recrystallized and the left half was not.

Figure 9: Scans of polished and Sopori-etched samples, recrystallized under the gentle profile.

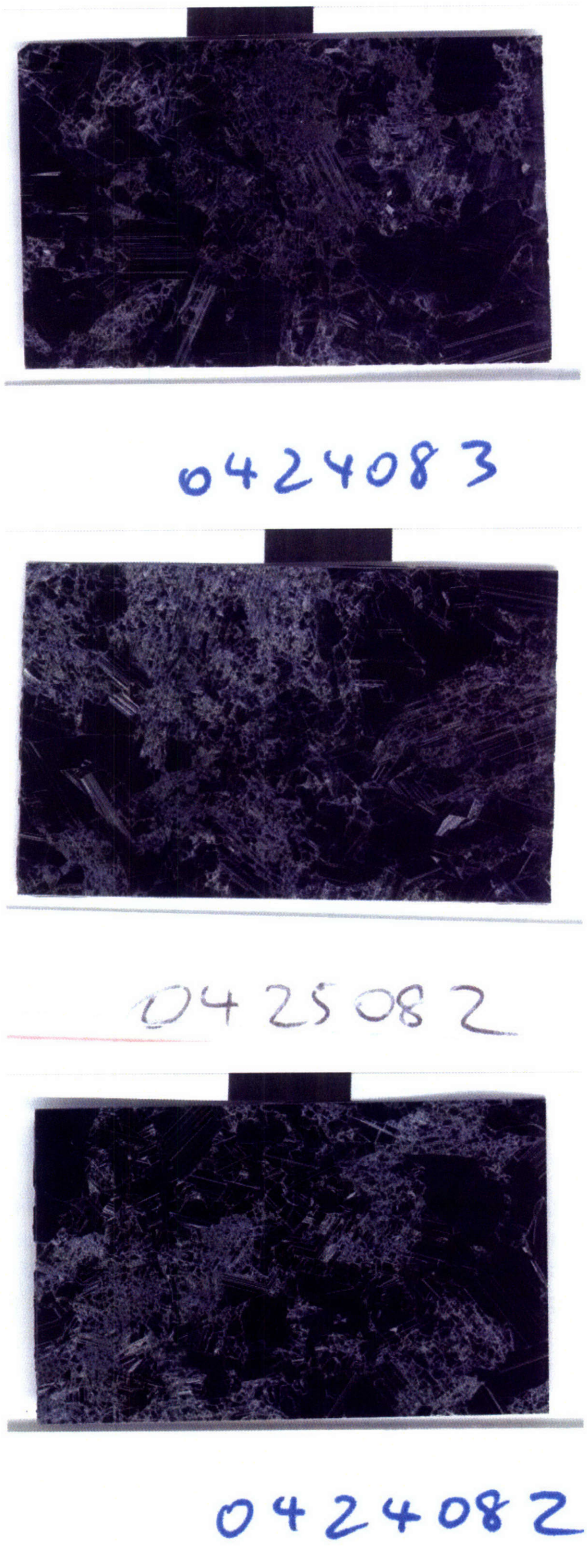

Figure 10: Scans of Sopori-etched control samples. They were not recrystallized, only damage etched.

Histograms of Dislocation Density

In scanned samples, pixel brightness distribution can be represented by binning pixel brightness values into 128 bins and plotting the population of each bin.

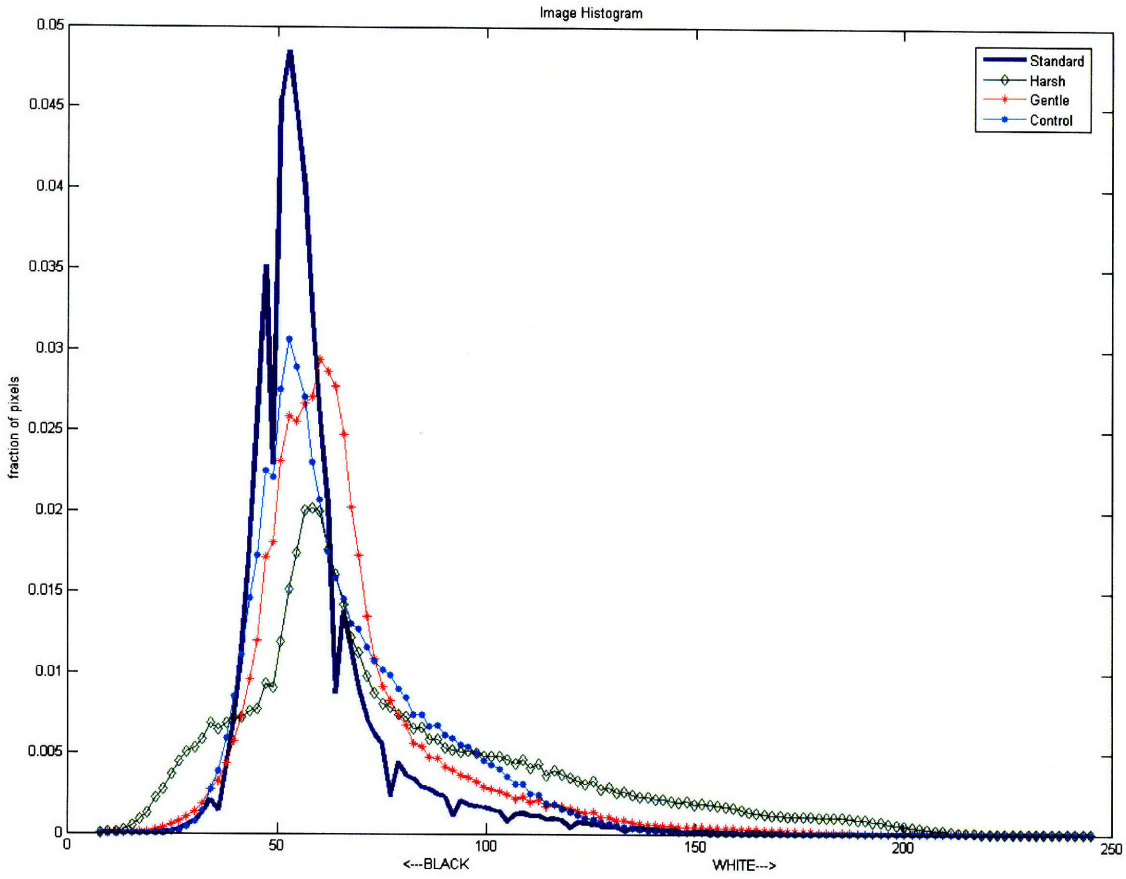

As is apparent from the histogram, unrecrystallized samples are higher in dislocation density than samples recrystallized under the standard profile. While bulk dislocation densities are comparable between the unrecrystallized (control) and standard-recrystallized samples (apparent in peaks at the same brightness level), a much larger

proportion of the recrystallized samples' surface area was at the (low) bulk dislocation density level.

Thus, overall dislocation density in the recrystallized samples improved. This improvement is visible in the tail of the distribution: the control samples exhibit higher pixel density than standard-recrystallized samples in the range above 65 on the histogram, corresponding to more wafer area at dislocation densities above $1E6cm^{-2}$.

Peak dislocation density of standard, gentle, and control samples are comparable, as apparent in the convergence of their tails on the bright end. The harsh profile has a clearly higher peak dislocation density than the other three types; the tail terminates at a brightness value over 200.

Among all recrystallized samples, standard profile samples had the lowest bulk dislocation density, followed by gentle and harsh profiles (which had similar bulk densities).

6.3 Minority Carrier Lifetime

Sample Preparation

In order to achieve accurate lifetime measurements, it was necessary to passivate the surface of wafers under test. Without passivation, measured lifetimes reflect surface-limited recombination as opposed to bulk-limited recombination, which is not informative about the overall finished material quality. To passivate, samples were etched for 3 min in 10:1 hydrofluoric acid and then sealed in a plastic bag containing a mixture of 0.08 molar iodine in methanol.

Figure 12: Sample bagged in iodine-methanol solution. All samples were measured for minority carrier lifetime in this type of packaging to minimize surface limited recombinbation.

Quasi Steady-State Photoconductance

Testing on the Sinton WCT-120 quasi steady-state photoconductance lifetime tester at an injection level of 1E15 cm^{-3} showed a minority carrier lifetime of 27μs in the longest-lifetime unrecrystallized sample and 20μs in the longest-lifetime recrystallized sample, which was recrystallized under a harsh temperature profile.

Minority carrier densities in recrystallized samples at the measured lifetimes were low. The cause of this low minority carrier density remains under investigation. Lifetime results for a representative sample from each profile are as follows:

Profile	**Sample Designation**	**Minority Carrier Density (cm^{-3})**	**Best Lifetime (μs)**
Control	0424083	1E15	27
Standard	0306087	1E14	3
Harsh	0407082	1E15	20
Gentle	0407084	2E14	6

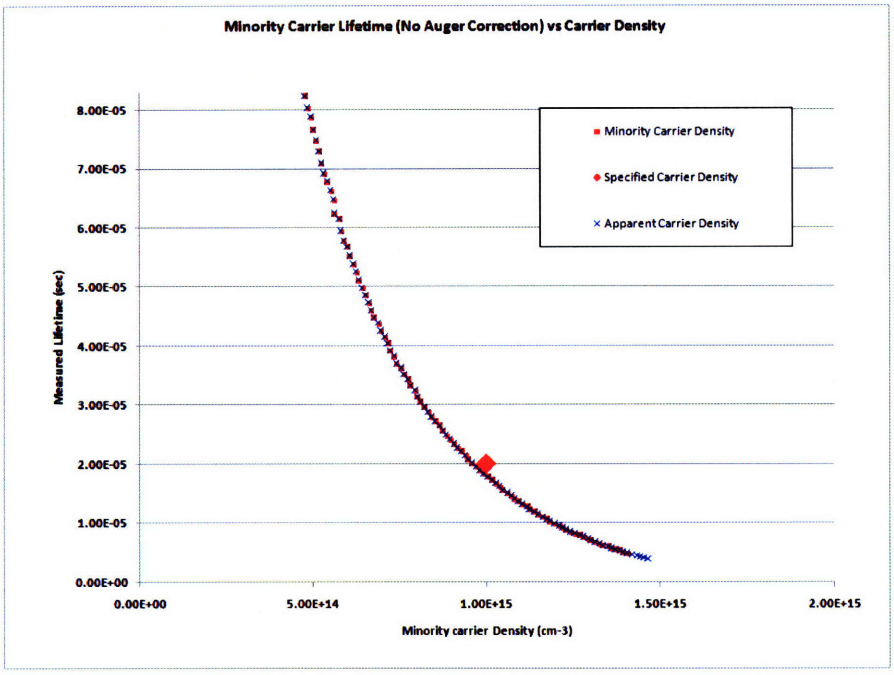

(a)

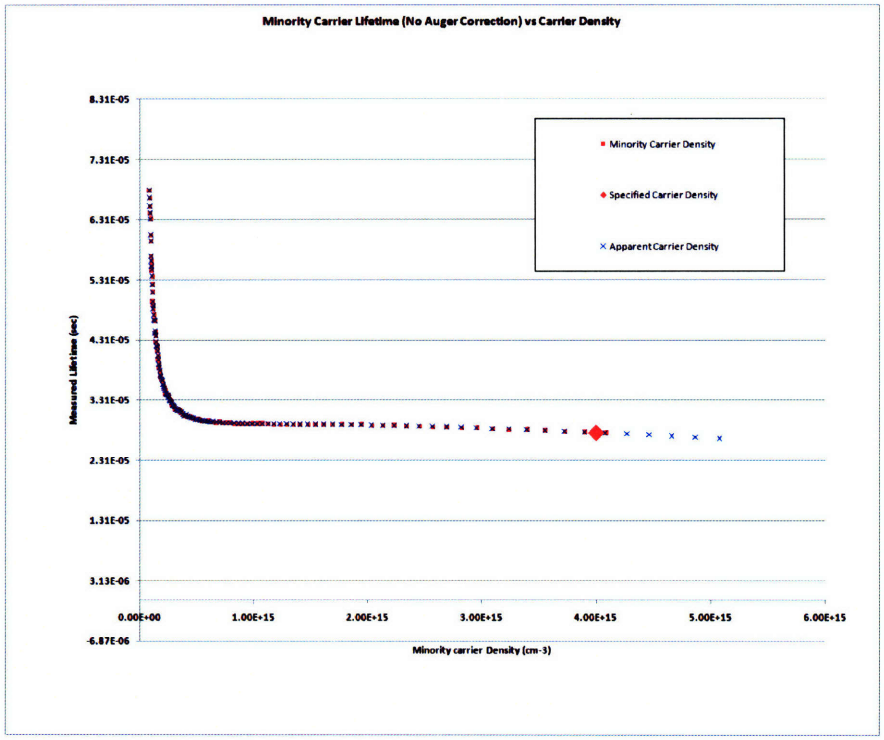

(b)

Figure 13: Lifetime vs. minority carrier density in (a) the best recrystallized sample – from a harsh profile – and in (b) the best control sample. There is a notable difference in the slope of the plots; this likely derives from the lower overall photoconductance of the recrystallized sample.

Semilab Lifetime Mapping

Probably owing to the low photoconductance of recrystallized samples, the μ-PCD process employed at Semilab's Burlington, MA facility was unable to resolve a signal on recrystallized samples. The cause for this marked decrease in photoconductance likely relates to the presence of impurities in recrystallized samples, as dislocation density is lower in recrystallized samples than in control samples. Further analysis, such as ICP-MS, could be employed in the future to determine whether impurities are indeed the cause of decreased photoconductance.

4. Partial Melting Anomalies

Because of reliability issues, several recrystallization trials were terminated prematurely when heater element contacts burned out. As a result, some samples were melted only partially. These samples provide insight into changes in dislocation density, grain structure, and the shape of the melt front.

Figure 14: This is a flatbed scan of a polished and dislocation-etched wafer. Recrystallized material is on the left; unrecrystallized material is on the right. This sample was being recrystallized under a harsh temperature profile when heater power was lost, midway through the trial. Dislocation density decreases in a fairly smooth gradient across the melt front into the recrystallized material.

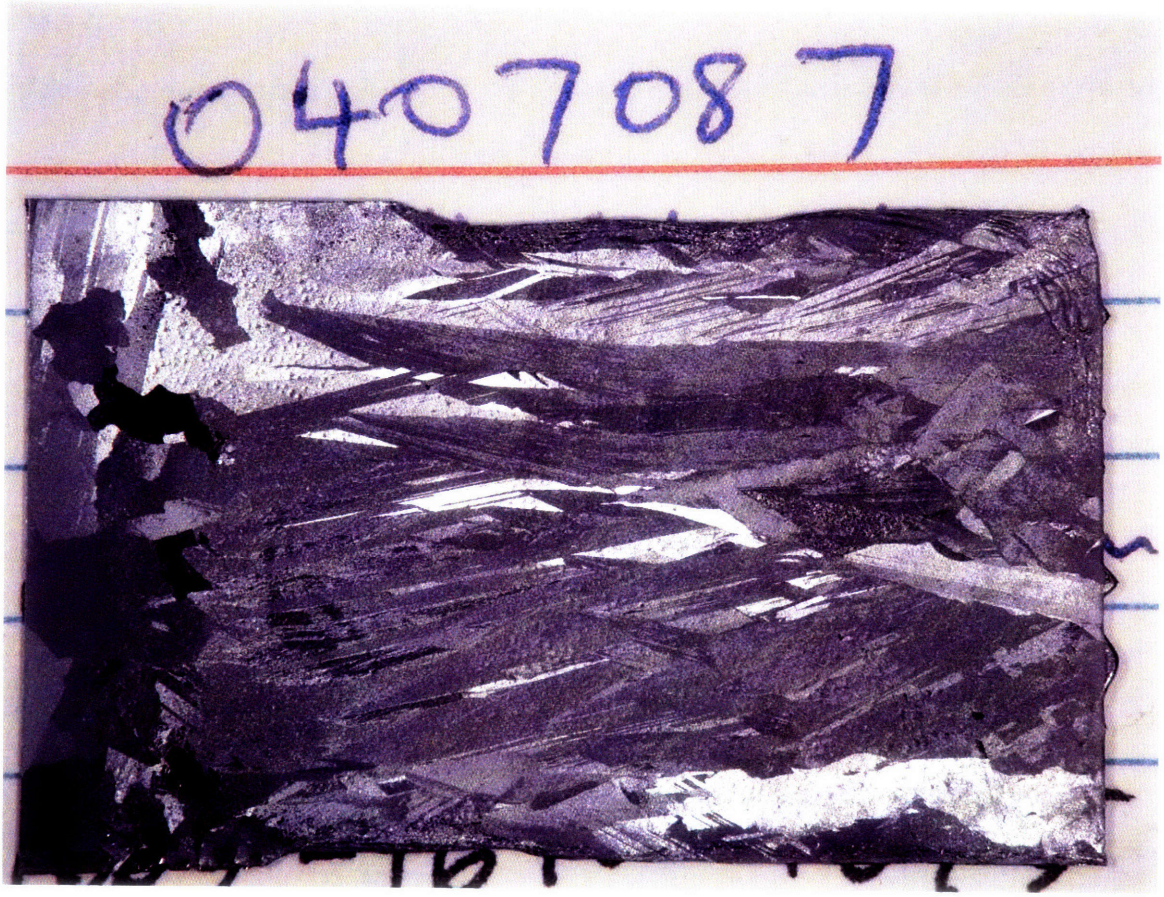

Figure 15: Camera image of a texture-etched sample, in which the heater power was lost most of the way through. The sample was being recrystallized under a gentle temperature profile. Recrystallized material is on the right; unrecrystallized material is on the left. Differences in principal dimensions of the grains are apparent.

Chapter 7: Discussion and Conclusion

7.1 Interpreting the Results

Grain Size

As explained in Gilman, grains allowed to equilibrate in a strained solid tend to be smaller in size than grains equilibrated in an unstrained solid.[1] Strain (which – in this case – arises from temperature gradients) scales as the second partial derivative of temperature profile in any given direction, as explained in 2.2. Thus, in examining the three temperature profiles used for recrystallization, the extent of downward concavity of the profile just after melting is the quantity of interest. High temperature ranges are also of interest, as dislocations are most mobile and susceptible to propagation just after freezing.

Although the macroscopic shapes of the harsh and gentle temperature profiles are smoother, the standard profile is nearly linear in the range between 100mm and 175mm from the furnace entrance. Conversely, the harsh and gentle profiles both undertake a rapid temperature rolloff in this critical range of positions. They are more strongly concave down, which translates to higher stress. Especially since the stress is experienced at a high temperature, this difference in temperature profile is likely a large contributing factor to the differences in grain size. Since samples under the standard profile are subject to nearly linear temperature profiles just after freezing, the relative coarseness of their grains makes sense.

Dislocation Density

Dislocation formation is strongly related to thermoelastic strain during crystal growth, as explained in 2.2. This agrees with the model of temperature profile concavity

put forth above. Because of linearity between 100mm to 175mm from the door, the standard sample was under relatively little thermal stress during the crucial period just after freezing. Accordingly, it displayed lower dislocation density than the harsh and gentle samples. These results are consistent with what is apparent in the histogram of 6.2; indeed, the harsh profile (which is the most severely concave down at its peak) yielded by far the highest dislocation density. These results agree with the temperature profile theory explained in 2.2.

Lifetime

Minority carrier lifetime results are somewhat anomalous. Specifically, the sample with the highest measured dislocation density yielded the longest lifetime. This should not be the case, as dislocations function as recombination centers for charge carriers.

Absent further data, a plausible explanation for this behavior is that some samples were subject to contamination at an unknown point in the process while other samples were spared the contamination. ICP-MS could be used to investigate the purity of low-lifetime samples, in an attempt to pinpoint the source of recombination.

2. Conclusions

Viability of Process

While dislocation density results were promising and lifetime results were ambiguous (though not disastrous), the actual viability of this process will not be known until recrystallized wafers are processed into working solar cells. More relevantly, the viability will not be known until a microcrystalline wafer (as opposed to a

multicrystalline wafer) is processed into a functioning device, to replicate the true grain structure of a high-speed cast wafer.

To be able to carry out those tests, the first hurdle will be recrystallizing a wafer with sufficiently predictable geometry to be handled in all downstream cell fab processes (all of which are designed to work with flat, smooth wafers). The rough surface finish and non-uniform perimeter may pose processing problems.

In that vein, the commercial viability of this process hinges (among other factors) on the ability to output geometrically consistent wafers. If the process does not succeed by that metric, it is no better than conventional wire sawing. Thus, it is critical in this case that wafer geometry is not destroyed by recrystallization. Cast cells may well be at a performance disadvantage to conventionally-manufactured cells, so they need every cost advantage possible – in this case, the cost advantage is eliminating wire sawing.

Another factor driving the commercial viability of this process is the hardware capital and maintenance costs. As with any prototype, the recrystallization hardware set was unreliable and demanding of significant operator attention. How readily the hardware scales has yet to be determined.

3. Future Work

Temperature Profiles

Although these results are consistent with the temperature profile theory put forth in Chapter 2, the actual temperature profiles yielded were not ideal for studying different temperature profile concavities. These tests could be carried out again with more highly-

optimized temperature profiles, perhaps using a more sophisticated thermal model of the furnace or different heater setpoints.

That the standard profile improved dislocation density was (at least to some extent) simply a matter of luck; a different temperature profile might have yielded far worse or far better dislocation density. Improved modeling and thermal mapping will allow more informed decisions about how to set up an optimized temperature profile.

Lifetime and Photoconductance

Because of the ambiguous data from QSSPC lifetime measurement and the non-existent date from the μ-PCD test, further effort should be applied to investigating the drop in photoconductance of recrystallized samples, as well as taking a more comprehensive set of lifetime measurements. Once the causes are understood, focus can shift to producing high-lifetime wafers repeatably.

Starting Wafers

As described above, processing a microcrystalline wafer into a multicrystalline wafer will be an immensely valuable means of verifying the process for its real intended end use. Other variations on starting wafers should be attempted as well: float zone samples polished on both sides (to investigate nucleation mechanisms), as well as exceedingly thin (<150μm) samples (to simulate industry-standard wafer thicknesses).

Capsule

Several options beyond thermal oxide exist for encapsulating the wafer before recrystallization. Si_3N_4, wet oxide, spun-on glass coatings, or even a self-applied oxide (in the recrystallization furnace) might offer advantages over the existing standard capsule – not the least of which might be invested energy and Takt time.

Geometric Constraint

As described above, the ability to produce a dimensionally repeatable wafer is essential to scaling the process. Much remains to be understood about the mass transfer mechanisms at work during recrystallization, as well as about the tendencies of samples to deviate from rectilinearity.

Purity Analysis

To determine the extent of impurity migration into recrystallized samples, tests such as EDX, ICP-MS, or SIMS should be carried out on starting and recrystallized wafers. Once the relative contributions of various contaminants are known, the process of minimizing their impact (in the interest of increasing minority carrier lifetime) can begin. This will become increasingly important as the hardware is scaled up and new materials find their way into the furnace.

References

---

[1] Burgers., W. G. "Principles of Recrystallization." *The Art and Science of Growing Crystals*. Ed. Gilman. New York: Wiley and Sons, 1963. 421.