High-Temperature Defect Engineering for Silicon Solar Cells: Predictive Process Simulation and Synchrotron-Based Microcharacterization

by

David P. Fenning

M.S., Mechanical Engineering
Massachusetts Institute of Technology (2010)
B.S., Mechanical Engineering
Leland Stanford Junior University (2008)

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

Author .....
Department of Mechanical Engineering
May 23, 2013

Certified by .....
Tonio Buonassisi
Associate Professor of Mechanical Engineering
Thesis Supervisor

Accepted by .....
David Hardt, Professor of Mechanical Engineering
Chairman, Department Committee on Graduate Theses

High-Temperature Defect Engineering for Silicon Solar Cells: Predictive Process Simulation and Synchrotron-Based Microcharacterization

by

David P. Fenning

Submitted to the Department of Mechanical Engineering
on May 23, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

Efficiency is a major lever for cost reduction in crystalline silicon solar cells, which dominate the photovoltaics market but cannot yet compete subsidy-free in most areas. Iron impurities are a key performance-limiting defect present in commercial and pre-commercial silicon solar cell materials, affecting devices at concentrations below even one part per billion. The lack of process simulation tools that account for the behavior of such impurities hinders efforts at increasing efficiency in commercial materials and slows the time-to-market for novel materials.

To address the need for predictive process modeling focused on the impact of impurities, the Impurity-to-Efficiency kinetics simulation tool is developed to predict solar cell efficiency from initial iron contamination levels. The modeling effort focuses on iron because it is known to limit most industrial solar cells. The simulation models phosphorus diffusion, the coupled diffusion and segregation of iron to the high phosphorus concentration emitter, and the dissolution and growth of iron-silicide precipitates. The 1D process simulation can be solved in about 1 minute assuming standard processing conditions, allowing for rapid iteration. By wrapping the kinetics simulation tool with a genetic algorithm, global optima in the high-dimensional processing parameter space can be pursued for a given starting metal concentration and distribution.

To inform and test the model, synchrotron-based X-ray fluorescence is employed with beam spot sizes less than 200 nm to identify iron-rich precipitates down to 10 nm in radius in industrial and research materials. Experimental X-ray fluorescence data confirm model predictions that iron remains in heavily-contaminated multicrystalline materials after a typical industrial phosphorus diffusion. Similar measurements of the iron-silicide precipitate distribution in multicrystalline silicon samples before and after higher-temperature gettering steps confirm that the higher the process temperature, the larger the reduction in precipitated iron, leading to marked lifetime improvement.

By combining the impurity kinetics modeling with the experimental assessment

of metal distribution, design guidelines for process improvement are proposed: the high-temperature portion of the process can be designed to enhance dissolution of precipitated iron, while the cooldown from the high-temperature process is crucial to the reduction of the interstitial iron concentration.

Finally, while precipitated iron reduction improves with higher temperatures, some regions of multicrystalline silicon samples degrade with higher-temperature gettering steps. To investigate the effect of gettering temperature on the remaining lifetime-limiting defects, spatially-resolved lifetime, interstitial iron concentration, and dislocation density are measured. The detailed defect characterization and analysis provide insight into the limitations of high-temperature phosphorus diffusion gettering.

Thesis Supervisor: Tonio Buonassisi
Title: Associate Professor of Mechanical Engineering

Acknowledgements

It is deeply humbling to reflect on the wealth of both personal and professional support I have received during my graduate research. Foremost, I would like to thank my advisor Professor Tonio Buonassisi for opening his door to me, for his energy, and for his vision. His mentorship over the past five years on how to form important research questions and develop strong experimental approaches made this dissertation possible. The ideas and support of my committee members, Prof. Sam Allen and Prof. Jeff Grossman, were critical to strengthening my thesis and their broader guidance tremendously appreciated.

For much of my learning on the ground, in the lab, I have the group postdocs to thank, especially Jasmin Hofstetter and Mariana Bertoni – two great scientists who know how to get things done. Not only would the synchrotron runs have been impossible without their help during crunch time, but their insights and the examples they set were truly formative to my development as a scientist. A big thank you is in order as well for our lab research scientists, especially Alexandria Fecych and David Berney Needleman, for their helping hands and minds in and out of the lab. I would be remiss not to thank my fellow graduate students in the PV Lab, especially Joe Sullivan, Sergio Castellanos, Doug Powell, and Ashley Morishige not just for their intellectual contributions, but for their friendship throughout this journey. It has been truly an honor to work each day with such amazing colleagues.

Collaborators from around the world have helped me formulate and test the hypotheses developed in this work and contributed invaluable samples for study. A special thanks to Annika Zuschlag and Giso Hahn, Jean-François Lelièvre and Carlos del Cañizo, Gianluca Coletti, and Markus Rinio. I would also like to thank Barry Lai for his expertise and support for my work at 2-ID-D. To the many other grad students across the MechE department and across MIT who have helped me with their stimulating ideas, with hands-on help given in lab, and for their friendship, I thank you. The list of individuals deserving of acknowledgement extends well beyond these pages and includes the entire MIT and PV communities.

I also feel it important to acknowledge the critical support for my doctoral research in solar provided by the National Science Foundation's graduate research fellowship and the Martin Family Society of Fellows for Sustainability at MIT.

Finally, I would like to thank my family and dear friends for their unwavering, loving support and my wife, Molly, for her understanding and for her ability to challenge me to improve in all aspects of life.

Dedicated in memory of Joan H. Hill.

Contents

List of Figures

List of Tables

References to Published Work

The iron kinetics simulation described in Chapter 2 and used throughout this thesis appears in detail in:

> J. Hofstetter, D. P. Fenning, M. I. Bertoni, J. F. Lelièvre, C. del Cañizo, T. Buonassisi. "Impurity-to-Efficiency Simulator: Predictive Simulation of Silicon Solar Cell Performance Based on Iron Content and Distribution, *Prog. Photovolt.: Res. App.*, Vol. 19, Issue 4, 487-497 (2011).

Chapters 4, 5, and 6 have largely appeared in the following three articles:

> D. P. Fenning, J. Hofstetter, M. I. Bertoni, G. Coletti, B. Lai, C. del Cañizo, T. Buonassisi. "Precipitated Iron: A Limit on Gettering Efficacy in Multicrystalline Silicon," *Journal of Applied Physics*, **113**, 044251 (2013).
>
> D. P. Fenning, J. Hofstetter, M. I. Bertoni, S. Hudelson, M. Rinio, J. F. Lelièvre, B. Lai, C. del Cañizo, and T. Buonassisi. "Iron Distribution in Silicon After Solar Cell Processing: Synchrotron Analysis and Predictive Modeling," *Applied Physics Letters*, **98**, 162103 (2011);
>
> D. P. Fenning*, Annika Zuschlag*, M. I. Bertoni, B. Lai, G. Hahn, T. Buonassisi. "Improved Iron Gettering in Contaminated Multicrystalline Silicon by High-Temperature Phosphorus Diffusion," *Journal of Applied Physics*, (accepted). *equal contributions.

The XRF data on metal-dislocation interactions of Chapter 8 has been published as:

M. I. Bertoni, D. P. Fenning, M. Rinio, J. Maser, T. Buonassisi. "Nanoprobe X-ray Fluorescence Characterization of Defects in Large-Area Solar Cells," *Energy and Environmental Science*, **4**, pp. 4252-4257, (2011).

The work contained in this thesis has also been presented at:

D. P. Fenning, A. S. Zuschlag, A. Frey, M. I. Bertoni, B. Lai, G. Hahn, T. Buonassisi, "Investigation of Lifetime-Limiting Defects after High-Temperature Phosphorus Diffusion in Silicon Solar Cell Materials," 39th IEEE Photovoltaics Specialist Conference, Tampa, FL, June 2013 (oral presentation).

D. P. Fenning, J. Hofstetter, A. Zuschlag, A. E. Morishige, G. Hahn, T. Buonassisi, "Iron Kinetics SImulation and Experimentation Demonstrating Potential for Novel Industrial Processing of Silicon Solar Cells," MRS Fall Meeting, Boston, MA November 2012 (oral presentation).

D. P. Fenning, M. I. Bertoni, V. Rose, J. Maser, T. Buonassisi, "Multiscale Synchrotron-based Study of the Physics and Device Impacts of Metal-Dislocation Interactions in Silicon Solar Cell Materials," MRS Fall Meeting, Boston, MA, November 2012 (oral presentation).

D. P. Fenning, J. Hofstetter, D. M. Powell, A. E. Morishige, A. Zuschlag, G. Hahn, & T. Buonassisi. "Design Guidelines for Tailored Phosphorus Diffusion Gettering," 22nd NREL Workshop on Crystalline Silicon Solar Cells & Modules, Vail, CO, July 2012 (invited oral presentation).

D. P. Fenning, M. I. Bertoni, T. Buonassisi. "Advancements in Synchrotron-Based Characteriation of Metals in Silicon," 5th International Workshop on crystalline Silicon Solar Cells, Boston, MA, November 2011 (poster presentation).

D. P. Fenning, J. Hofstetter, T. Buonassisi, "Optimizing Industrial Phosphorus Diffusion for Silicon Solar Cells", Gettering and Defect Engineering

in Semiconductor Technology XIV, Loipersdorf, Austria, September 2011 (poster presentation).

D. P. Fenning, T. Buonassisi. "Effective Iron Gettering in Lightly-Doped Emitters," 37th IEEE Photovoltaics Specialist Conference, Seattle, WA, June 2011 (oral presentation).

D. P. Fenning, J. Hofstetter, M I. Bertoni, J. F. Lelièvre, C. del Cañizo, T. Buonassisi. "Synchrotron-based Microanalysis of Iron Distribution after Thermal Processing and Predictive Modeling of Resulting Solar Cell Efficiency," 35th IEEE Photovoltaics Specialist Conference, Honolulu, HI, June 2010 (oral presentation).

Chapter 1

A Brief Introduction to Silicon Materials for Photovoltaics

The potential for silicon photovoltaics to meet current and future energy demand is real. The solar resource, or the amount of power hitting the Earth's surface from the sun, is over 4,500x the power humans currently require [1]. The goal of the photovoltaics industry is to collect just a fraction of that incoming solar radiation and convert it directly into electricity for immediate use, for local storage, or for supplying power to the electrical grid. Despite tremendous growth in recent years, including 81.5% year on year growth in 2011 [2], the PV industry must significantly reduce the cost of solar power to reach grid parity, where the cost of solar power is less than or equal to the cost of power produced by the varied, often carbon-heavy sources supplying the grid [1].

Within photovoltaic technologies, devices based on crystalline silicon continue to dominate the market with a market share of about 87% as of 2011 [3]. In some ways, development of silicon photovoltaics is a byproduct of knowledge of the silicon material system stemming from the success of the semiconductor industry, where silicon has been studied and engineered for over fifty years [4]. On the other hand, silicon is a natural choice for photovoltaic devices. The bandgap of silicon is ≈ 1.1 eV, nearly optimal for a single-junction photovoltaic device [5]. Silicon is also the second most abundant element in the Earth's crust, meaning that there should be no

shortage in supply of raw materials even when silicon solar cells are produced on the terawatt scale.

While silicon-bearing raw materials are widely available, a significant barrier for silicon solar cell manufacturers is the cost of highly-purified elemental silicon. In fact, the purified feedstock is the largest single component of the cost of producing a silicon PV module, equaling approximately 15% of the current module cost [6]. The feedstock is expensive because purity requirements are stringent. Semiconductor

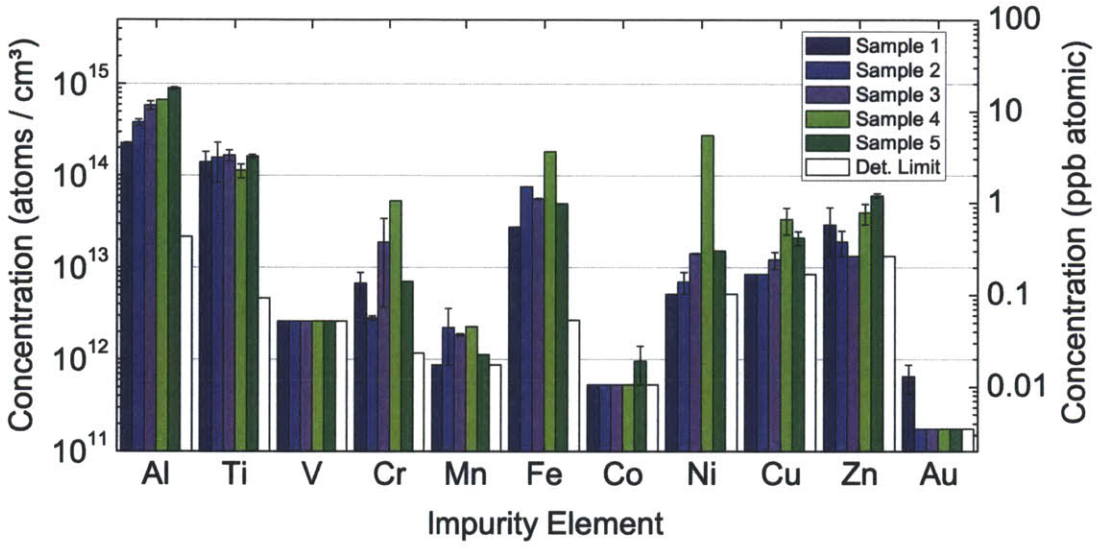

Figure 1-1: Inductively-coupled mass spectrometry measurements are shown for five multicrystalline sample taken from the upper portion of a middle brick in an industrially cast ingot. A wide range of impurities are detected, including many of the deleterious *3d* metals, and a number of metals reach concentrations above one part per billion.

materials and devices are highly sensitive to dilute concentrations of metal impurities [4, 7]. Silicon-based photovoltaic devices are no exception, with conversion efficiencies typically decreasing as metal impurities exceed atomic concentrations of parts per billion [8]. Sensitive mass spectrometry methods, such as inductively-coupled plasma mass spectrometry (ICPMS), can be capable of analyzing the bulk concentrations of most metals in silicon with better than per billion detection limits. Impurity concentrations measured by ICPMS for five multicrystalline silicon samples from the upper portion of a typical cast ingot are shown in Fig. 1-1. As can be seen, a number

of impurity elements reach part per billion concentrations in these wafers. Such concentrations are equivalent to a few of drops of water in an Olympic size swimming pool.

Reducing the impact of impurities on crystalline silicon solar cell performance could reduce costs in several ways. If device efficiency could be maintained despite the presence of higher levels of impurities, it would enable the recycling of tops, sides, and tails of ingots, and even the wire-sawn kerf. Improved material usage could have a drastic impact on silicon feedstock cost and, consequently, solar cell manufacturing cost. Long lead times for new high-purity polysilicon plants also exacerbate feedstock price fluctuations associated with overcapacity and undercapacity. Alternative routes (*e.g.* upgraded metallurgical grade silicon) offer faster ramping times and lower costs [9], but have been plagued by high impurity contents.

Iron impurities in particular limit the bulk minority carrier lifetime of most as-grown p-type silicon wafers [10, 11]. Iron is particularly deleterious in p-type silicon in part because of its large electron capture cross section [12, 13], but also because of its inevitable incorporation in wafers due to its high diffusivity and solubility in silicon [14–16]. Furthermore, iron is readily present in feedstocks [17], crystal growth crucibles and their linings [18], and throughout the industrial device manufacturing environment [4, 19].

Recently, several authors have intentionally iron-contaminated their starting materials to isolate the macroscopic device-level efficiency impacts of iron contamination [20–23] updating foundational Westinghouse studies of metal contamination in silicon solar cells [8] to assess the impact of advances in substrate material quality, higher-efficiency cell architectures, and improved processing. Intentional contamination in the melt produces selected impurity concentrations above "background" impurity levels like those seen in the standard mc-Si wafers in Fig. 1-1, allowing for the specific study of the selected contaminant while maintaining the same thermal history for the impurities as would be seen after a standard crystal growth. These benchmark studies have helped researchers update the specifications of tolerance limits of metal impurities in a standard silicon solar cell process [24, 25]. The defect concentration

limits extracted (e.g., 25 ppma Fe in the feedstock) are similar to the landmark 1980 Westinghouse study of impurities in silicon solar cells [8], in large part because even as solar cell processing and architectures have advanced, so have efficiency targets. The end result is the same despite improved processing: even in ingots cast using electronic-grade silicon, a large fraction of solar cells, specifically from the top, bottom, and sides of the ingot, have unacceptably low efficiencies due in part to high metal concentrations that exceed the specified limits [21, 22].

The top of the ingot is degraded because a phenomenon known as solid-to-liquid segregation occurs during growth. The thermodynamic solubility of metals in solid silicon is significantly lower than in liquid metal-silicon solutions, often expressed as a segregation coefficient at the melting temperature, k_{m} , which is on the order of $\approx 10^{-4} - 10^{-6}$ for most metals in silicon [14, 26]. As the solidification front moves from the bottom of the crucible toward the top during the directional solidification of the ingot, the segregation coefficient means there exists a large driving force to keep a large fraction of the total number of metal impurities in an ever-shrinking liquid fraction. As a result, the metal concentrations in the liquid layer grow rapidly as the liquid fraction shrinks. Toward the end of solidification, the remaining heavily-contaminated liquid at the top of the ingot solidifies ultimately resulting in an 10-1000x accumulation of metal concentrations in the solidified crystal as solidification nears completion. This solid-to-liquid segegration is well described when assuming a well-mixed liquid and using simple mass conservation, producing the Scheil distribution:

$$C(x) = k_{\mathrm{m}} C_o (1 - x)^{k_{\mathrm{m}} - 1} \tag{1.1}$$

where C is the concentration of impurity at the ingot height, x, expressed as a fraction of the total ingot height, and C_o is the original concentration of impurity in the melt.

In addition to the tops of ingots, wafers from the red zone near the crucible wall typically exceed allowable impurity concentrations [27]. Due to the high temperatures during solidification ($T^{\mathrm{Si}}_{\mathrm{melting}} = 1414°\mathrm{C}$), impurity in-diffusion of fast-diffusing metal impurities, *i.e.* Cu, Ni, Fe and Cr [14], from the crucible lining causes a "red zone" of

poor quality material around the edges, top, and bottom of the ingot. Iron, because of its large capture cross section for electrons and high solid solubility, has been shown to cause significant degradation of material quality in border regions of p-type ingots, leaving lifetimes often below 1 μs [27]. This region of degraded materials affects a considerable fraction of all wafers produced by such methods [26].

The entirety of the poor-quality red zone is visualized in a minority carrier lifetime map of a cast multicrystalline ingot cross-section in Fig. 1-2. The lifetime measure-

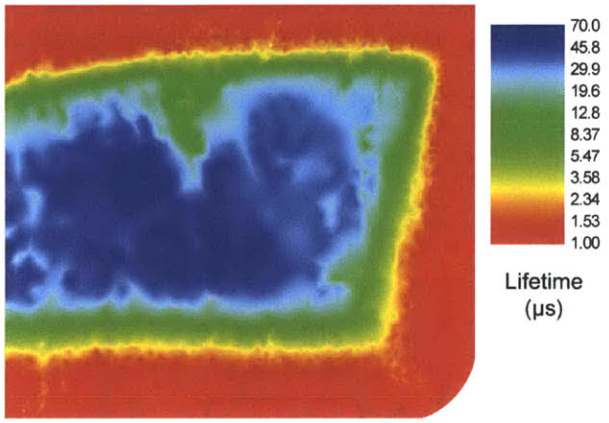

Figure 1-2: A lifetime map of a vertical cross-section of half of a cast mc-Si ingot clearly reveals the red zone at the top, bottom, and sides of the ingot caused by the influence of metal impurities. Reproduced from [28].

ment shown in Fig. 1-2 was done by carrier density imaging [29]. The red zone areas seen are readily predicted by modeling the solid state diffusion of iron from the walls of the crucible, solid-to-liquid segregation of iron during the growth process using an effective segregation coefficient to account for the diffusion boundary layer that forms due to finite solidification speed, and back diffusion of iron after solidification [28].

As a result of this contamination during crystal growth, the tops, bottoms, and sides of an ingot are simply discarded after crystal growth in a traditional silicon solar cell process. When these losses are added to the material lost during wafering to sawdust (known as kerf loss), the end result is a material utilization of generally less than 50% of the original high-purity silicon feedstock [9]. This loss of pure, crystalline silicon obtained through a cost- and energy-intensive process adds significantly to the

subsequent cost of solar power [30, 31]. At the same time, there is a strong push in the industry to higher efficiencies, making a significant fraction of cells produced today that end up in the bottom tail of the distribution unacceptable because of high impurity concentrations [32]. One path to achieve tighter solar cell efficiency distributions with higher average values, while also achieving better silicon utilization, is to recover high carrier lifetimes in previously discarded wafers or wafers that would otherwise turn into unsalable cells.

To improve the material quality of as-grown wafers cut from such ingots and achieve highly efficient devices, the metal contamination must be carefully managed during processing [33]. A schematic of the time-temperature profile of the solar cell manufacturing process is shown in Fig. 1-3. After cooldown from crystallization, the

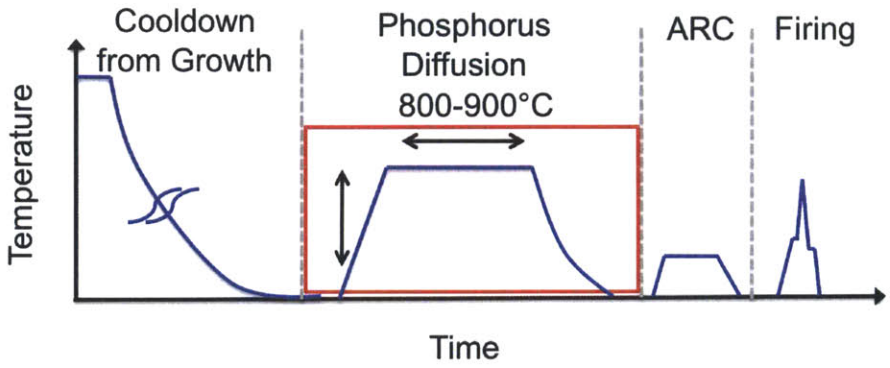

Figure 1-3: After sawing wafers from cast ingots, the first high-temperature process in solar cell manufacturing is phosphorus diffusion, typically conducted between 800-900°C. After phosphorus diffusion, subsequent elevated temperatures are normally seen during the deposition of an anti-reflection coating (ARC, $T < 450$°C) and finally during metallization firing, which typically includes a temperature spike above 800°C for several seconds. Phosphorus diffusion is both the highest-temperature and the longest step in the manufacturing sequence that converts silicon wafers to solar cells.

wafers are sawn from the ingot, cleaned, and textured to enhance light capturing. The first high-temperature step of the cell process is phosphorus diffusion to form the n-type emitter layer. Carried out typically at 800-900°C, phosphorus diffusion is the highest-temperature and longest step of the cell process. Afterward, an anti-reflection coating (ARC) is deposited, typically a silicon nitride layer deposited via

plasma-enhanced chemical vapor deposition at temperatures below 450°C. Finally, the front and back sides are metallized before a firing step is carried out to form Ohmic contacts with a temperature spike to around 800°C for at most several seconds.

Silicon solar cell performance depends strongly on how well the cell fabrication process described above is matched to the quality of the wafer source material, because the impact of metal contamination depends on its chemical state and spatial distribution, which in turn depend on processing conditions [33–37]. Because of the non-linear, coupled effects of cell processing on the electronic properties of the material, the final charge carrier lifetime of a solar cell cannot be easily deduced from the initial lifetime of the wafer [38]. Given the breadth of material quality in the crystalline silicon market – from high purity, single crystal (sc-Si) grown by the Czochralski technique using electronic grade Si (up to nine "nines" purity, or 99.9999999%) to directionally-solidified upgraded metallurgical grade multicrystalline (mc-Si) wafers – one might expect an equally varied set of manufacturing conditions. However, such source-specific tailoring of cell processing is limited by the high-dimensional parameter space of cell processing and the resulting monetary and temporal cost associated with conducting a design of experiment to find efficiency maxima. To date, the photovoltaic industry has required costly, high-purity silicon because technologies have not been developed to adequately mitigate the negative impact of metal impurities.

In particular, the time-temperature profiles in a standard solar cell process have been developed for electronic grade Si (EG-Si) feedstocks but are not well-adapted for solar grade Si (SoG-Si) because of the much higher defect concentrations in SoG-Si [24]. Due to the low total concentration of impurity atoms in EG-Si, metal impurities precipitated in the bulk dissolve completely at process temperatures, facilitating impurity removal of fast-diffusing species by means of external gettering during phosphorus emitter diffusion. In contrast, because of the high contamination levels in SoG-Si, high temperature steps have been predicted to lead to only the partial dissolution of metal precipitates during the short phosphorous diffusion gettering (PDG) applied in a typical industrial solar cell process [39, 40]. To extract the full efficiency potential from low-quality wafers, the industrial process needs to be adapted to ac-

commodate SoG-Si, as well as to enhance performance of cells originating from areas of the ingot that contain higher impurity contents, notably edge, top and bottom regions. Thus far, the flexible adaptation of the solar cell fabrication process to each source material has been a resource-intensive experimental task due to the lack of predictive modeling capability.

This thesis is centered around the development of a predictive model of how solar cell processing determines the iron impurity distribution. The Impurities-to-Efficiency (I2E) simulation described herein incorporates only what we determine to be the essential physical mechanisms of iron redistribution to streamline the integration of the coupled partial differential equations and minimize obfuscation of the driving factors. Using synchrotron-based micro-characterization, this thesis provides measurement of the distribution of metals at the start of solar cell process for a variety of industrially-relevant materials. The impact of solar cell processing on this metal distribution is then evaluated, both by simulation and experimental synchrotron-based measurement, to asses the effect of both standard processing and several alternative processing concepts. Ultimately, in the latter chapters, design guidelines are extracted from these experimental and simulation results that should provide insight and guidance for effective control of iron and the tailoring of the solar cell process to the starting material.

1.1 Thesis Overview

The goal of this body of research is to establish design guidelines for advanced high-temperature processing of low-cost wafer materials to achieve high-efficiency solar cells. To accomplish this objective, the development of accurate impurity gettering models is pursued in tandem with synchrotron-based microscopic investigations of precipitated metal distribution as a function of processing conditions.

In Chapter 2, I briefly review the physical mechanisms phosphorus diffusion gettering of iron and how researchers in the field have developed models to describe the gettering process. I then describe the components and function of the impurity-to-

efficiency model that I co-developed with Dr. Jasmin Hofstetter to predict solar cell performance from initial iron contamination levels.

In Chapter 3, I detail the methods used in this thesis to measure interstitial and precipitated iron. In particular, I describe the use of synchrotron-based X-ray fluorescence microprobes (μ-XRF) at Argonne National Laboratory's Advanced Photon Source to measure the iron precipitate distribution in as-grown and processed wafers. With beam spot sizes less than 200 nm, we are able to identify iron precipitates smaller than 10 nm in radius [41] in a wide range of silicon solar cell materials.

With these theoretical and experimental foundations in hand, in Chapter 4 I explore the response to a typical industrial solar cell process of laboratory-grown research ingots that were intentionally contaminated with iron during crystal growth using both the I2E simulation tool and synchrotron-based X-ray fluorescence. The experiments provide measurements of the microscopic as-grown iron precipitate distribution used as inputs for the model. The μ-XRF data allow me to produce the first large area (100s μm^2) quantification of the microscopic iron distribution in multicrystalline silicon ingots as a function of position and processing condition. Additionally, these experiments provide a baseline for understanding the impact of phosphorus diffusion gettering. Following the model predictions, the experimental data confirm that iron remains in heavily contaminated materials after the high-temperature processing typical of an industrial solar cell process, demonstrating that the use of highly-contaminated materials in industry demands a more effective impurity reduction during solar cell processing to maintain solar cell efficiencies.

Looking to improve upon baseline processing, in Chapter 5and 6 I investigate two alternative processing schemes. In Chapter 5, I measure how an already-established alternative process called low-temperature annealing, which has been successfully applied to achieve higher solar cell efficiencies with heavily iron-contaminated materials, changes the iron distribution. Simulations of low-temperature annealing match experimental XRF results well, and the results add insight into the mechanisms behind lifetime improvement during low-temperature annealing [42].

In Chapter 6, I focus on the high-temperature plateau portion of the solar cell

process. By elevating the gettering temperature, increased precipitate dissolution is seen by μ-XRF measurements of the same nanoscale precipitates before and after processing. The increased reduction in precipitated metals is found to be strongly correlated with improved lifetimes in heavily-contaminated materials.

Collecting the contributions of this thesis, a broad design approach for process improvement is presented in Chapter 7. In short summary, the high-temperature portion of the emitter formation process can be designed depending on initial total iron concentration to enhance dissolution of precipitated iron, while the cooldown from the high-temperature process determines the segregation gettering and reprecipitation of dissolved iron. Together the high-temperature step and its cooldown determine the final iron distribution, and the tradeoff between process time and lifetime improvement from iron redistribution can be made through optimization studies.

Lastly, the limitations of high-temperature defect engineering are explored in Chapter 8. Results illustrating the impact of dislocations and dislocation-metal impurity interactions on the efficacy of the phosphorus diffusion gettering process are detailed.

Results and conclusions are summarized in Chapter 9, before I offer several directions and open questions regarding high-temperature defect engineering of silicon solar cells.

Chapter 2

The Physical Mechanisms of Phosphorus Diffusion Gettering and their Modeling

> Our model of Nature should not be like a building – a handsome structure for the populace to admire, until in the course of time some one takes away a corner stone and the edifice comes toppling down. It should be like an engine with movable parts. We need not fix the position of any one lever; that is to be adjusted from time to time as the latest observations indicate. The aim of the theorist is to know the train of wheels which the lever sets in motion – that binding of the parts which is the soul of the engine.
>
> Sir Arthur Stanley Eddington

In photovoltaic devices, the material is fundamentally the device – there are no moving parts that convert mechanical motion to electrical energy. That is not to say that the material itself can be considered static. During the high-temperature processing involved in solar cell manufacturing, the material undergoes significant transformations that change it from a passive raw semiconductor into a functional energy material (cite solar). Solid-state defect reactions and their kinetics typically limit the solar cells produced by dictating the transformation of the raw material, and the success of large-scale, low-cost production of devices hinges upon a rigorous optimization of this material transformation from passive crystal to active electronic device. What is more, we must accomplish this difficult process-structure-property optimization challenge as quickly as possible to allow the rapid deployment of low-carbon renewable energy technologies.

In this chapter, I will give an overview of the approach taken in this thesis to simulate the response of silicon wafer materials to solar cell processing. To predict the impact of processing on downstream device performance, focus is placed on simulating device-limiting iron impurities and their redistribution within the crystal structure during solar processing, specifically during the longest, highest temperature step of the process where the largest changes in iron distribution are seen: phosphorus diffusion (cf. FIg. 1-3) . First, I begin by describing the general distribution of metal impurities in as-grown silicon materials. I then describe the mechanisms by which phosphorus diffusion affects this metal distribution. Having reviewed the complexities of phosphorus diffusion, I describe several approaches by other researchers to model the impact of phosphorus diffusion. I then present our impurity-to-efficiency (I2E) model and the assumptions therein which are used to produce a streamlined, engineering approach to solving for the distribution and impact of iron after phosphorus diffusion.

2.1 The Distribution and Device Impacts of Metal Impurities

In general, there are three classes of metal impurity defects in crystalline silicon: (1) atomic point defects where the metal atom is dissolved in solid silicon solution at low concentrations [43], either interstitially or at substitional points on the silicon lattice (2) metal-silicide second-phase precipitates that occur at much lower spatial density but are 10s of nanometers in size [34], and (3) inclusions from the manufacturing environment that enter the melt and often contain metal-oxides that are 100s or 1000s or nanometers in size but have little electronic impact and so are not focused on here. Fig. 2-1 illustrates the typical distribution of metals within silicon wafers.

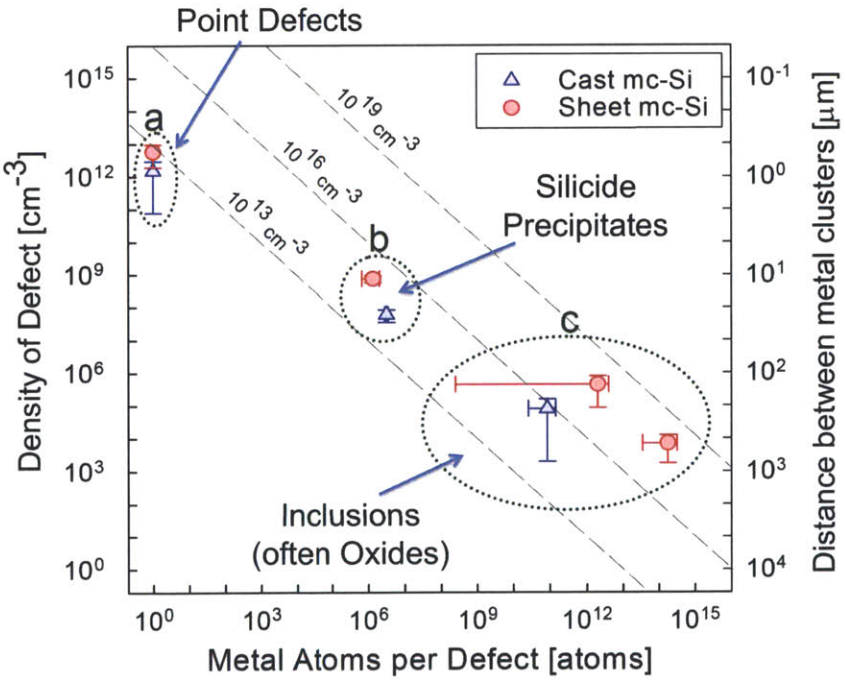

Figure 2-1: The density of impurity defects vs the number of atoms per defect is shown. Synchrotron-based X-ray fluorescence and absorption spectroscopy have enabled the direct observation of silicide and oxide precipitates in industry standard materials. Point defects are unobservable at synchrotron microprobes, but their high spatial density means that they strongly affect the electronic performance of the cell. Total metal concentrations are shown on the diagonals. Figure adapted from [34].

2.1.1 Recombination at metal point defects

The principal electronic effect of metal impurities in semiconductors is to enhance the recombination of electron-hole pairs by creating deep level, midgap energy states when in their point defect chemical state [7, 44]. Point defects are also the most problematic for electronic devices because they occur in the highest spatial densities, in effect short-circuiting charge flow through the device. Recombination at deep levels created by point defects is well described by the Shockley-Read-Hall model [45, 46] for recombination at impurity centers, shown schematically in Fig. 2-2. The lifetime

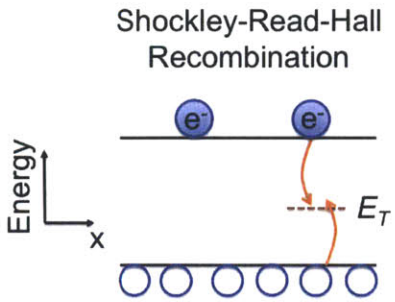

Figure 2-2: Shockley-Read-Hall recombination at a defect trap state in p-type silicon is depicted on a band diagram. The model describes recombination at midgap trap states introduced by impurity defects. The midgap states provide efficient recombination centers for electron-hole pairs, greatly decreasing the minority carrier lifetime.

dictated by recombination at impurity centers according to simplified Shockley-Read-Hall (SRH) statistics is:

$$\tau_{SRH} = \frac{\tau_{p_0}\left(n_0 + n_1 + \Delta n\right) + \tau_{n_0}\left(p_0 + p_1 + \Delta p\right)}{n_0 + p_0 + \Delta n} \tag{2.1}$$

which applies for many experimental conditions in silicon solar materials (with non-degenerate, steady-state conditions with recombination center density much less than the background doping level) [45, 47]. According to the model, the time constant for electron capture, τ_{n_0}, is governed by the density of defects in the material, N, the thermal velocity of electrons and holes, ν_{th}, and the capture cross-section of the

defect, σ_n:

$$\tau_{n_0} = \frac{1}{N\sigma_n \nu_{th}} \tag{2.2}$$

In addition, the energy level of the defect trap state with respect to the valence or conduction band determines the population of electrons and holes when the Fermi level is at the defect level, according to:

$$n_1 = N_c \exp\left(\frac{E_T - E_C}{\mathrm{k}_B T}\right) \tag{2.3}$$

where N_c is the conduction band density of states, E_T is the energy level of the defect, E_C is the conduction band energy level, k_B is the Boltzmann constant, and T the temperature. Analogous terms can be determined for the hole capture time constant, τ_{p_0}, and the population of unfilled defect states, p_1. It is important to note the $1/N$ dependence of lifetime on defect density – high concentrations of metal impurity point defects have a hyperbolic impact on minority carrier lifetime.

2.1.2 Recombination at Precipitates

While precipitates do affect the minority carrier lifetime of the material as well [48, 49], their direct electronic impact is generally much less than point defects because of their lower densities. Several models exist that describe the direct impact of metal impurity precipitates on lifetime.

del Cañizo *et al.* [48] developed a model that describes recombination at metal precipitates solely in terms of minority carrier kinetics, including no field effect. By solving the continuity equation for minority carriers in the presence of precipitates that have a recombination active surface, the model determines the recombination rate attributable to precipitates. Precipitates of characteristic radius, r_0, and density, N_{pcp}, provide surfaces for recombination dictated by carrier diffusivity, D, and precipitate surface recombination velocity, s, according to:

$$\frac{1}{\tau_{pcp}} = 4\pi {r_0}^2 N_{pcp} \frac{sD/r_0}{s + D/r_0} \tag{2.4}$$

An alternative approach was taken by Plekhanov and Tan [50], who modeled the capture cross-section of precipitates attributing carrier recombination to thermionic emission across the Schottky junction formed at the precipitate-Si interface rather than to recombination at the surface of the precipitate. Their model resulted in capture cross sections that are generally 10-100x the geometric cross-section, with field effects becoming more important the smaller the precipitates.

Experimental measurement of the characteristic capture cross-section of precipitated metals in silicon has been difficult to achieve because of the need for high spatial resolution (< 1 μm) in both chemical mapping of the precipitate composition and the electrical characterization of precipitate recombination activity. Deep level transient spectroscopy, though highly successful at characterizing capture cross-sections of point defects, is much less useful for precipitate characterization because the amplitudes of broad spectral peaks caused by deep states at precipitates no longer correspond to the concentration of defects [51]. Recently, Gundel *et al.* [52] used high-injection, confocal micro-photoluminescence imaging [53] in correlation with micro-Raman mapping to identify stress as a contributing factor in recombination activity at precipitates. However, the relative contributions of stress, precipitate size, and Schottky-diode effects in determining recombination activity at precipitates remain an area of active investigation.

In addition to the direct recombination of minority carriers at precipitates, the presence of precipitates at pre-breakdown sites suggests silicide precipitates may cause detrimental shunting of the junction [54, 55]. Further, Chapter 4 will demonstrate that the dissolution of precipitated metals in response to solar cell processing largely determines the final concentration of interstitial impurities, which in turn is of critical importance to the final minority carrier lifetime.

The final effective minority carrier lifetime of the material, τ_{eff}, which can be correlated with the final efficiency of the device [38], is the harmonic sum of the minority carrier lifetimes dictated by each type of recombination in the material. For example, lifetime can be limited due to recombination at point defects, at precipitates, at the surface of the sample [56], due to Auger recombination at high injection [57, 58],

etc.,

$$\frac{1}{\tau_{eff}} = \frac{1}{\tau_{M_i}} + \frac{1}{\tau_{M_{pcp}}} + \frac{1}{\tau_{surf}} + \frac{1}{\tau_{Auger}} .... \tag{2.5}$$

2.2 Gettering: What is it?

To decrease the density and concentration of impurities and improve the quality of semiconductor materials, researchers in semiconductor and photovoltaic technology have developed a suite of process knowledge and annealing techniques collectively known as defect engineering. For silicon photovoltaics, defect engineering involves mainly the manipulation of metal impurities, according to processes collectively known as gettering. Gettering steps include processes aimed at the removal of metal impurities from the bulk of the device and/or the reduction of the defect density. At its most basic, a gettering process involves the release of a metal impurity from a bulk site, the diffusion of the impurity to a gettering region, and the capture of that impurity in the gettering region, as shown in Fig. 2-3.

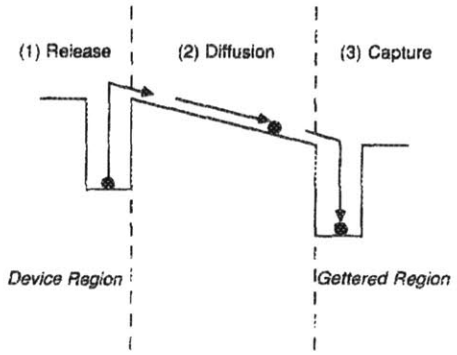

Figure 2-3: A cartoon of a gettering process shows the basic steps involved. First, an impurity becomes mobile in the bulk, after which it diffuses to a gettering region, where it remains, effectively "captured" by the gettering region. Figure from [59].

Gettering can occur by a number of different mechanisms which fall into two main categories: relaxation gettering and segregation gettering. Relaxation gettering occurs when supersaturated interstitial metal impurities precipitate out at favorable heterogeneous nucleation sites. Internal relaxation gettering occurs within the bulk,

either nucleating at or re-precipitating at existing metal precipitates at structural defects such as grain boundaries and dislocations [60, 61] or oxygen precipitates [62, 63]. Internal gettering is one of the principle tools used to produce exceptionally high-quality active regions near the surface of the device in Czochralski monocrystalline wafers for the semiconductor industry, wherein controlled oxygen precipitation [64] is used to internally getter metal impurities [15]. External relaxation gettering can occur when precipitation is favored at surfaces rather than within the bulk, which occurs for example in the case of Cu in monocrystalline silicon [65]. Relaxation gettering depends on supersaturation and so generally occurs only during cooling.

Segregation gettering occurs when the impurity species has a higher solubility in one device region than in the bulk and so has a driving force for diffusion into the gettering layer: a mechanism that operates even during annealing at high temperatures. For example, during Al back contact firing the low eutectic temperature of the Al-Si system at 577°Cresults in a liquid layer forming which has much higher solubility for metals than the semiconductor bulk [66, 67]. Segregation gettering can also rely on the solubility enhancement due to the Fermi-level dependent solubility of metal impurities which have charged defect states in the bandgap [14]. For example, Gilles *et al.* [68] found significantly enhanced solubility for many transition metals, including iron, in both n^{++} and p^{++} silicon. The pairing of charged metal impurities with ionized dopant atoms and an increase in the number of substitutional impurities can also contribute to enhanced solubility and lead to segregation gettering [4].

2.3 Phosphorus Diffusion Gettering

Phosphorus diffusion gettering (PDG) is a special case involving multiple equilibrium and non-equilibrium gettering mechanisms. Overall, phosphorus diffusion gettering can be highly effective at removing metal impurities from the bulk wafer, often leading to 1-2 orders of magnitude reduction of impurity concentrations including iron in both monocrystalline and multicrystalline silicon [69, 70].

The gettering efficiency for both interstitial and precipitation iron reduction is

limited by the solid solubility and the diffusivity of interstitial iron (Fe_i) in Si, both of which have an Arrhenius relationship with temperature. Consequently, the dissolution of $FeSi_2$ precipitates increases with increasing temperatures. Phosphorus diffusion gettering thus represents a delicate balance: on one hand, iron-silicide precipitates must be dissolved to reduce bulk metal content; on the other, residual bulk iron point defect concentrations must be small to improve minority carrier lifetimes. The

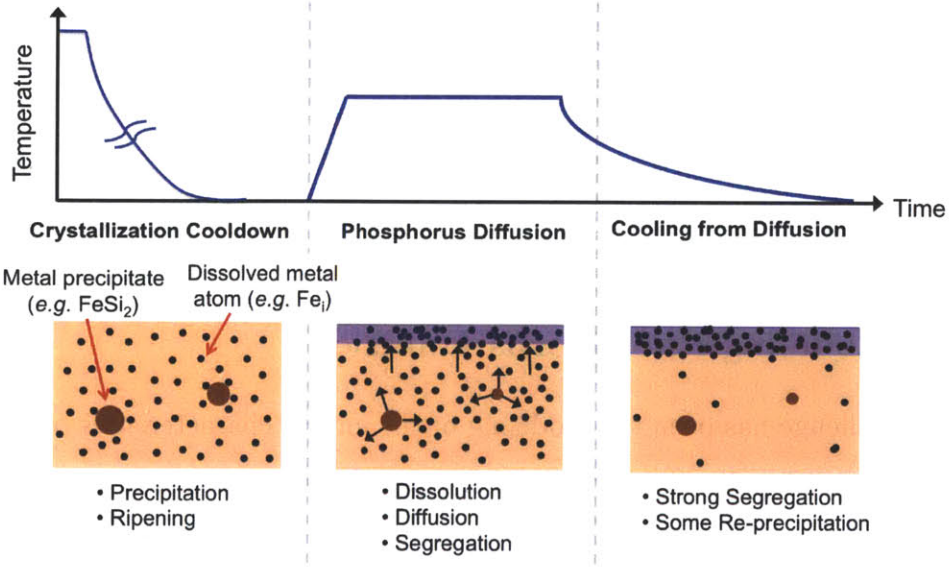

Figure 2-4: After crystallization, a heterogeneous distribution of metal impurities is found, with much of the total concentration contained in metal-silicide precipitates in equilibrium with metal point defects in solution. The high temperatures of phosphorus diffusion cause the partial dissolution of the metal-silicide precipitates. The n^+-layer that forms from the in-diffusing phosphorus acts as a sink for many of the resulting interstitial metal impurities. Strong segregation gettering of impurities to the emitter layer is possible during the cooldown from phosphorus diffusion The trade-off between the dissolution of precipitates injecting interstitial metals into the bulk and their sinking at the gettering layer at the surface is critical to the final device performance.

exploration of this trade-off has traditionally been conducted in a trial-and-error fashion, guided by intuition gained from the integrated-circuit industry. Examples of the multiple ongoing defect reactions during phosphorus diffusion are depicted in Fig. 2-4. In some cases, phosphorus diffusion is observed to result in decreased performance [59, 71, 72]. Additionally, the effectiveness of a given P-diffusion process

is observed to be highly dependent on the starting material [73–75], complicating efforts to develop universal annealing profiles. The development of a model to describe the most important features of the phosphorus diffusion gettering of lifetime-limiting iron impurities is detailed below.

2.4 Modeling Phosphorus Diffusion Gettering

To model phosphorus diffusion gettering of iron, we require applicable physics in several areas: phosphorus diffusion, iron point defect diffusion, coupling impurity kinetics to phosphorus diffusion iron, and the behavior of second-phase iron precipitates during gettering.

2.4.1 Modeling Phosphorus Diffusion

The first challenge has been the modeling of the unique characteristics of phosphorus diffusion. Traditionally the formation of the p-n junction during silicon solar cell manufacturing is carried out by the thermal diffusion of phosphorus into a boron-doped p-type wafer to form a 0.5-1 μm thick n^+-layer at the surface [76]. Typically, the surface source of phosphorus is formed either by chemical vapor deposition from the gas phase or sprayed or spun-on in liquid phase. For PV, phosphorus oxychloride, $POCl_3$, has become the dominant species for phosphorus delivery. Picked up by bubbling an inert carrier gas, typically N_2, through a liquid $POCl_3$ source, the phosphorus oxychloride oxidizes at process temperature to form a phosphosilicate glass on the surface of the wafer. The phosphorus glass is reduced at the silicon interface, forming interfacial silicon dioxide and yielding elemental phosphorus which diffuses into the wafer. The reaction proceeds as [77]:

$$4\,POCl_3 + 3\,O_2 \longrightarrow 2\,P_2O_5 + 6\,Cl_2$$
$$2\,P_2O_5 + 5\,Si \longrightarrow 4\,P + 5\,SiO_2$$

Using a $POCl_3$ source for P-diffusion has the additional benefit of producing Cl_2, which acts as a strong chemical getter for metals at the wafer, boat, or furnace tube

surfaces by reacting with them to form volatile compounds and helping maintain a clean diffusion environment [4].

The p-n junction forms around where the phosphorus concentration meets the background boron doping concentration to create a space-charge region. The fixed space charge of the ionized dopants generates an associated electric field that separates excited carriers and allows the energy of carriers excited by solar radiation to be collected. Because of the high effective activation energy of diffusion for phosphorous (2.1 ± 0.1 eV), the diffusion is often conducted at temperatures above 800°C [78]. The solid solubility of phosphorus in silicon is above 10^{20} atoms/cm^3 at temperatures above 800°C [79, 80], far greater than a typical background boron doping of $\approx 10^{16}$ atoms/cm^3, and so the junction depth is controlled by the diffusion length of phosphorus for a given time-temperature profile.

The diffusion of phosphorus results in a characteristic "kink-and-tail" profile that shows strong departures from a simple complimentary error function diffusion equation solution. Instead, the diffusion profile can roughly be considered the sum of two diffusion more typical Gaussian diffusion overlaying one another. Accordingly, building off the groundwork of Fair and Tsai [81], Hu *et al.* [82] proposed a two-mechanism model by which a vacancy diffusion mechanism was assigned to the slower-diffusing, high-concentration component of the phosphorus profile, while an intersticialcy mechanism was use to explain the faster-diffusing, low-concentration component. Collecting ideas from these and other previous works, Uematsu [83] built an integrated diffusion model to to take into account the vacancy mechanism, kick-out, and the Frank-Turnbull mechanism for not only phosphorus but also boron and arsenic diffusion in silicon. Recently, Bentzen *et al.* [78] added a Boltzmann-Matano analysis to deconstruct the contributions of each mechanism to the concentration-dependent phosphours diffuvisity, allowing for accurate, semi-empirical modeling of phosphorus diffusion profiles over a range of time-temperature profiles of revelance to silicon solar cells at concentrations up to the electrically-active limit of $\approx 3 \times 10^{20}$ atoms/cm^3, beyond which phosphorus clustering can account for the behavior [84]. Besides the time-temperature profile of the diffusion process, Bentzen *et al.*'s model requires only

one fit parameter – the concentration of phosphorus at the surface.

In an effort to remove even this single parameter, several authors (see *e.g.*, [85, 86]) have worked to establish predictive models of the oxide and glass growth at the surface, which in turn determines the phosphorus boundary condition for the diffusion problem. This work has yielded important kinetic relations such as the parabolic growth of the surface oxide with time and the rate-limiting step of phosphorus diffusion across the oxide. Ultimately, the goal would be to predict the phosphorus diffusion profile only from the gas ratios ($POCl_3$:O_2:N_2), but the uniqueness of each furnace geometry and configuration leave the chemical and fluid dynamic simulation of the chemical vapor deposition process out of reach.

More recently, interest in ion-implantation of phosphorus for advanced cell architectures such as selective emitter designs has increased [87]. Even for ion implantation, a high temperature anneal after implantation has been shown to provide performance benefits [88]. Regardless of the technique used to introduce phosphorus, the high-temperature anneal associated with the junction formation offers the opportunity for redistribution of many *3d* metal impurity defects because of their high solid-state solubility and diffusivities in silicon, providing the release and rapid diffusion required for gettering.

2.4.2 Coupling Impurity Behavior to Phosphorus Diffusion

The near-surface, highly-doped n-type layer acts as a sink for impurities, in part because the solubility of many impurities increases as the Fermi level moves toward the conduction band in the n-type layer [68], yielding another instance of segregation gettering. Some species, iron in particular, also form complexes with phosphorus that may increase the effective solubility [89, 90]. While there is little evidence for substitutional iron in n-type silicon [16], well-known substitutional impurities, such as Au and Pt, are also influenced by the non-equilibrium injection of silicon self-interstitials that is coupled to phosphorus diffusion [91–93], in a process known as injection gettering where the self-interstitials kick-out substitutional impurities leading to an accelerated gettering.

For the *3d* elements that are primarily or exclusively interstitial in silicon, such as Co and Fe, no significant injection gettering effect can be expected. However, segregation gettering alone cannot fully explain the magnitude of many phosphorus diffusion gettering results [94]. One possibility is that higher-temperature diffusions have been observed to result in SiP precipitation at the Si/P_2O_5 interface, which may encourage the precipitation of metal impurities locally and increase the power of the sink layer [95]. Further detailed reviews of the mechanisms of phosphorus diffusion gettering can be found in Meyers *et al.*[94], Istratov *et al.* [4], Schöter *et al.* [96], and Seibt *et al.* [97]. The mechanisms that dominate the capture of impurities in the gettering layer that develops during phosphorus are still under debate, providing an obstacle to modeling phosphorus diffusion gettering and its device impacts.

To overcome this position, Haarahiltunen *et al.* [98] developed a semi-empirical model of the segregation of iron to the phosphorus-rich layer by fitting SIMS data for phosphorus diffusion profiles and the correlated iron profile. The Haarahiltunen model for the segregation coefficient of iron relies on two chemical rate constants to describe the final concentration profile of iron in phosphorus-diffusion gettered wafers. While not required for the application of the model, Haarahiltunen hypothesizes that the two physical reactions that enhance solubility in n-type silicon include the substitution of Fe on Si lattice sites and the pairing of substitutional Fe with phosphorus [98]. The activation energies in the Haarahiltunen segregation model were updated including lower-temperature data in [99].

In contrast to the semi-empirical approach, other authors have made efforts toward a complete description of defect reactions involved in phosphorus diffusion gettering. Dunham *et al.* [100], for example, develop a rigorous analysis involving 8 defect pairing reactions, but containing many fitting parameters, for phosphorus diffusion. After simplification, Dunham *et al.* consider a 5 reaction model. Recently, Tryznadlowski *et al.* added the coupling of Fe point defects to the phosphorus diffusion to model phosphorus diffusion gettering of iron [101, 102], adding several more kinetic fitting parameters to their model, but unfortunately not providing the values used in their modeling.

2.5 Development of a Gettering Model Designed for Rapid Iteration: Impurity-to-Efficiency

To address the need for predictive process modeling focused on the impact of impurities, I worked with Dr. Jasmin Hofstetter, a postdoctoral scholar in the PVLab, to build what we call the impurities-to-efficiency, or I2E, simulation tool to predict solar cell efficiency from initial iron contamination levels. We focused on developing a practical, fast simulation that could be used in an iterative fashion to speed process development.

The one-dimensional I2E simulation tool consists of three serial components that capture the stages of industrial PV device production, starting from silicon wafers and ending with the solar cell. The three components, which have been independently and jointly validated using published experimental results [103], include: (1) a material kinetics model that calculates changes in the distribution of impurity iron and phosphorus atoms during annealing; (2) an electronic model that predicts depth-dependent minority carrier lifetime based on iron distribution; and (3) a device simulator that predicts solar cell performance based on the the device architecture and the calculated minority carrier lifetime distribution throughout the wafer.

(1) Phosphorus and Iron Kinetics Simulation

The heart of our model is in the impurity kinetics modeling, where we connect the work of several other researchers in the field to build the model. We model the concentration profile of three species: phosphorus dopant atoms, P, iron interstitials, Fe_i, and precipitated iron, Fe_p.

To model phosphorus diffusion, we use Bentzen's integrated diffusion model [78], where the phosphorus diffusivity is concentration-dependent:

$$\frac{\partial P}{\partial t} = \frac{\partial}{\partial x}\left(D_P(P)\frac{\partial P}{\partial x}\right) \tag{2.6}$$

and the diffusivity has two contributing components: vacancy-assisted diffusion at

high phosphorus concentration and self-interstitial assisted diffusion at lower concentrations. As noted above, the single fitting parameter in the phosphorus diffusion model is the surface concentration of phosphorus, which is set at 3×10^{20} cm^{-3} for all simulations in this thesis, equivalent to the limit of electrically active phosphorus and resulting in good agreement with measured diffusion profiles [103].

We describe the segregation of interstitial iron to the phosphorus rich layer through the use of You *et al.*'s diffusion-segregation equation [104, 105]:

$$\frac{\partial C_i}{\partial t} = \frac{\partial}{\partial x}\left(D_{Fe}\frac{\partial C_i}{\partial x} - \frac{C_i}{\sigma}\frac{\partial \sigma}{\partial x}\right) \tag{2.7}$$

where C_i is the local concentration of interstitial iron, D_{Fe} is the diffusivity of iron from [14], and σ is the depth-dependent segregation coefficient. The effective segregation coefficient used to describe the sinking of iron to the phosphorus-rich emitter is that developed in Haarahiltunen's semi-empirical approach [98]:

$$\sigma(x) = 1 + V_{Si}K_{eq,1}\frac{n(x)}{n_i}(1 + K_{eq,2}[P^+(x)]) \tag{2.8}$$

where V_{Si} is the concentration of Si vacancies and $K_{eq,1}$ and $K_{eq,2}$ are equilibrium constants characterizing the reactions between Fe interstitials, Si vacancies, and P. The coupling of iron to the phosphorus diffusion occurs in Eqn. 2.8 via the concentration of positively charged P atoms, $[P^+(x)]$, and the electron concentration, $n(x)$, which depend on the evolution of the phosphorus profile with time during the diffusion.

In addition to point defects, second-phase impurity precipitates must be taken into account, especially in multicrystalline materials, but also to account for gettering effects (like at oxygen precipitates) in monocrystalline silicon. Hieslmair *et al.* [106] originally developed a gettering simulator for internal gettering at oxygen precipitates that includes iron precipitate dissolution and growth based on Ham's law [107]. We likewise use Ham's law[107]:

$$\frac{dC_i}{dt} = 4\pi N_p r D\left(C_i - C_{eq}\right) \tag{2.9}$$

where C_i is the concentration of interstitial iron, D is the temperature-dependent diffusivity, and C_{eq} is the equilibrium iron solubility taken from Murphy and Falster's recent study of iron solubility [108].

(2) Lifetime Model

The final lifetime of the material is calculated assuming it is wholly limited by iron impurities, both interstitial and precipitated. The final effective minority carrier lifetime of the material, τ_{eff} is then:

$$\frac{1}{\tau_{eff}} = \frac{1}{\tau_{Fe_i}} + \frac{1}{\tau_{Fe_p}} \tag{2.10}$$

where SRH statistics are used to calculate the interstitial-limited lifetime (Eq. 2.1) and del Cañizo's surface recombination model (Eq 2.4) is use to model the precipitate-limited lifetime.

The lifetime model is meant to indicate the bulk lifetime in the p-type base that can be expected after gettering, and so we include no surface-limited or Auger terms as in Eq. 2.5. It is important to note that the lifetime in the emitter is not strongly affected by the accumulation of impurities in the n-type layer [109], due in part perhaps to the asymmetric minority carrier capture cross-sections of impurities like iron [13] but mainly because of the very short lifetimes that are dominated by Auger recombination in the n++-emitter. The impurities collected in the n-type layer are also tightly bound [94], and the presence of a phosphorus-rich layer can protect the bulk from contamination during later high-temperature processing [110].

In general, the lower the metal impurity defect density in the wafer bulk, the better the performance of the solar cell. Thus, not only the total concentration of impurities but also their distribution and density in the material is critical for establishing their electronic impact on the device [34, 111].

(3) Device Simulation

With a lifetime calculated as a function of depth in the wafer dependent upon the iron concentration, the standard PC1D [112] is used to simulate solar cell device efficiency. Device simulations require the input of many architecture-dependent parameters, and so we are restricted to analyzing cases where device parameters are well-known. We have made the I2E simulation tool available online at pv-i2e.mit.edu [113].

2.6 Assumptions and Limitations of the Impurities-to-Efficiency Model

Each of these pieces involves clear simplifications. For example, as an alternative to Ham's law, which is essentially a mean-field approximation [97], the Fokker-Planck equation has been used to simulate changes in precipitate size and density by Haarahiltunen *et al.* [40]. Haarahiltunen *et al.* [114] also added the ability to model the nucleation of iron during cooldown from crystal growth to simulate the starting distribution of iron precipitates in as-grown wafers. Schön *et al.* [49] continued to build off this modeling base and added an accounting of heterogenous nucleation at dislocations and of enhanced solubility in the surrounding strain field. The resulting 2D model is capable of describing iron nucleation at structural defects, phosphorus diffusion, iron diffusion and segregation to the phosphorus layer including the effect of enhanced solubility at dislocations, and the growth and complete dissolution of precipitates [115].

A judicious choice of assumptions can reduce the need for this complexity and dramatically reduce computational expense. For example, Fig. 2-5 shows the predicted lifetime after phosphorus diffusion using the Fokker-Planck formulation to study the effect of the assumption of the size distributions of as-grown precipitates. Typically, when using the simpler Ham's law formulation for precipitate dissolution and growth, a uniform precipitate size is assumed. If the wrong assumption for precipitate size is made, a wildly differing lifetime can result as seen in Fig. 2-5 where an assumption

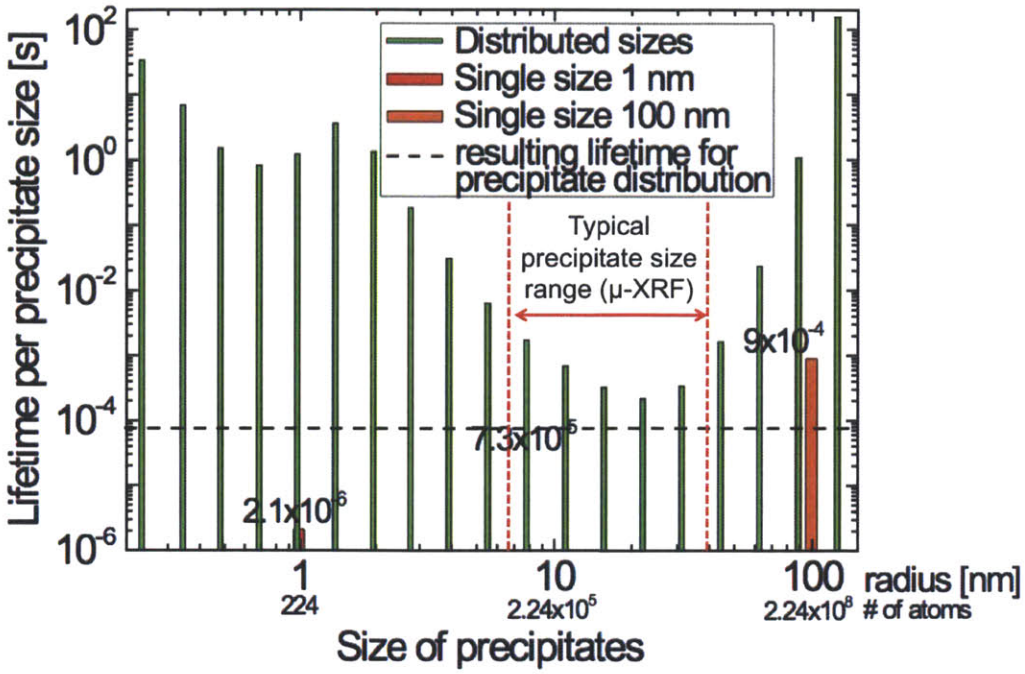

Figure 2-5: Shown here is the simulated lifetime after phosphorus diffusion as a function of precipitate size using the Fokker-Planck modeling formulation. Three different assumptions are made for the size of the as-grown precipitates in the calculations: (1) a varying precipitate size distribution from $<$ 1nm to $>$ 100nm (2) precipitates of all 1 nm in radius, and (3) precipitates of all 100 nm in radius. Clearly, the precipitate size assumption has a large impact on predicted final lifetime. Iron precipitates roughly 10-40 nm in size are seen to be lifetime-limiting in the case where there are a non-homogeneous distribution of precipitate sizes. Figure adapted from Schubert *et al.* [116]

of monodisperse 1 nm precipitates or 100 nm precipitates yields a final lifetime 2.5 orders of magnitude different. The power of the Fokker-Planck formulation is the ability to simulate a distribution of precipitate sizes and strive for more accurate results. It is important to note, though, that even in the distributed size simulation, precipitates between 10 and 40 nm in size almost completely dominate the lifetime after phosphorus diffusion, revealed by how close the overall lifetime at 73 μs (indicated by the horizontal dashed line) is to the lifetime limits impose by precipitates 10-40 nm in radius. Precipitates measured by μ-XRF later reveal precipitates predominantly in this size range. In fact, considering only the precipitates in the size

bin just larger than 20 nm in size would yield almost the same final lifetime. This simple result is because the final lifetime is the harmonic sum of the impact of all lifetime-limiting defects, and so the most pernicious of defects dominates the overall lifetime (e-quivalent to the weakest link in the chain). Fig. 2-5 demonstrates that much of the lifetime impact of iron precipitates can be collapsed onto precipitates of roughly 20 nm in size.

Clearly, the assumption of the precipitate size distribution plays a key role in modeling the effect of phosphorus diffusion gettering. Chapter 3 explains how synchrotron-based μ-X-ray-fluorescence is used in this thesis to measure the distribution of precipitate sizes to provide experimental evidence of what precipitate sizes are in relevant materials. Chapter 4 then describes an experiment involving measurement of the as-grown and phosphorus-diffused distributions of iron precipitates and reveals how a careful assumption of a monodisperse precipitate size distribution can provide good agreement with the experimental data.

The I2E model is also strongly dependent upon the assumption that iron is the lifetime limiting defect. While many as-grown multicrystalline wafers do appear limited by interstitial iron, after gettering the lifetimes can be limited by other defects. Furthermore, the heating and cooling rate of a prescribed process must fall within a domain of acceptable shear stresses in the wafer to avoid dislocation generation in the wafer. The generation of structural defects during high-temperature processing is not accounted for in the simulations, although under certain high temperature conditions it has been observed to cause lifetime degradation [72, 117]. Similarly, interactions of metals with dislocations is currently not included in the I2E simulation, although dislocation-dense regions are known to getter poorly with respect to low-dislocation areas [20, 118, 119]. In contrast to the large body of detailed experimental results of iron segregation to phosphorus layers, there is a need for further experimental exploration of the interaction of metals with dislocations to enable the extraction of the most important mechanisms. Several experiments studying the interaction of metals at dislocations and their impact on gettering and lifetime will be examined in Chapter 8.

By making these assumptions, the 1D I2E process simulation can be solved in about *1 minute* using standard processing conditions, allowing for rapid iteration and optimization in the high-dimensional process space. We published the framework of the model using experimental data from the literature to validate our model [103, 120], and the application of the model rather than its validation will be the focus of the thesis.

Constraints on Phosphorus Diffusion

There are also numerous constraints on a phosphorus diffusion process itself as it is a primary step in the production of the electronic device, and the goals of phosphorus diffusion in a multicrystalline silicon solar cell process are often contradictory. For high-purity silicon where impurity gettering may be of lower importance, an optimal diffusion process is one leading to a lightly-doped emitter, as lower dopant concentrations lead to lower emitter saturation currents, higher open circuit voltage, and improved blue response [121, 122].

In multicrystalline silicon, with its characteristically higher contamination levels [17], the optimization of gettering is more complex. Because impurity concentrations are generally higher than the limit of solid solubility at process temperatures [14, 24, 70, 123] , the gettering of impurities is kinetically limited. When impurity gettering is important to improve the bulk lifetime, a heavier diffusion may be favored. A thicker emitter, or at least one with a deep "tail", can also provide more effective protection from shunting during metallization firing, leading to higher fill factors [78].

In general, in this thesis I focus on purely the impact of iron gettering by phosphorus diffusion, without significant consideration for other device impacts. There are several reasons for this:

- By relaxing constraints, a broader parameter space can be explored in an effort to determine what is possible to achieve with lower quality materials, before applying industrial constraints.
- Even when considering the primary constraint of cost, the added cost of an extended phosphorus diffusion is small compared to the potential benefit in

contaminated materials. For example, recent cost modeling predicts that a doubling of process time can be offset by just a 1.5% improvement in *relative* efficiency, equivalent to a 0.25% improvement in absolute efficiency for a 17% industrial cell [124].

- The etch back of a strongly-diffused emitter has already been achieved for selective-emitter designs [125–127] and could be potentially be applied to produce uniformly better blue response for an industrial emitter after a heavy diffusion aimed at gettering

Overall, while absolute accuracy in a model is desirable, by streamlining the I2E model it is the goal of this thesis to extract gettering trends and guidelines rather than absolute values limited to particular materials. During phosphorus diffusion, the interplay between Fe_i formation by precipitate dissolution and the removal of Fe_i from the bulk (gettering) determines the final post-processed Fe_i concentration, which ultimately impacts solar cell performance. With the help of the I2E simulation tool [103, 128], we are able to evaluate these tradeoffs by calculating the post-processed precipitated iron concentration (Fe_p) and Fe_i concentration for nearly any given time-temperature profile and as-grown iron concentration and distribution.

Chapter 3

Quantitative Measurement and Analysis of the Distribution of Iron in Silicon

3.1 Introduction

The development of experimental techniques to characterize defects in silicon accurately at ppb detection limits has been critical to understanding defect reactions that govern device performance. Advances in mass spectroscopic techniques (e.g. GDMS, ICPMS), defect-pair dissociation (Fe-B), and synchrotron-based microfluorescence and spectroscopy have led to increasingly accurate understanding of the defects. In particular, of critical relevance to the modeling of phosphorus diffusion gettering of iron is the characterization of the iron distribution as a function of processing, including characterization of interstitial iron and precipitated iron.

Here, I describe the data collection and analysis use for quantifying the distribution of iron in silicon in this thesis. The goal is to establish a reliable procedure with minimal manual intervention for quantifying the data to maximize the accuracy and objectivity of the analysis. The key parameters to calculate here are the interstitial iron concentration, detailed in Section 3.2, and the size distribution and density of iron precipitates within the sample, discussed in Section 3.3.

3.2 Quantification of Iron Interstitials by Dissociation of Iron-Boron Pairs

In the case of points defects, electrical measurements can be sensitive to defect concentrations below part per trillion. This regime is well beyond mass spectroscopic methods, either bulk measurements like inductively-coupled plasma mass spectroscopy or local ion-beam techniques like secondary-ion mass spectroscopy (SIMS), and gamma-ray detection in neutron activation analysis. Deep level transient spectroscopy is a popular electrical characterization technique in the semiconductor defect field used to determine defect concentration, as the amplitude of the DLTS peak is proportional to the defect concentration [129, 130]. By varying the temperature and measuring the shift in the peak, the defect energy level can also be determined quite precisely, even in the presence of other defect levels. However, DLTS and many other electrical

techniques (*e.g.*, Hall effect) require metal contacts be applied to the sample. The additional processing required to deposit and anneal contacts to create Ohmic contacts for electrical measurement is undesirable because of the additional thermal steps, not to mention the potential for metal contamination of the bulk to occur.

3.2.1 Theory of Iron Point Defect Quantification

For point defects that form Coulombic pairs with ionized substitional dopants, Zoth and Bergholz [131] proposed a contactless detection method. Interstitial iron in particular pairs with boron acceptors at room temperature when samples are left in the dark, according to the reaction:

$$Fe_i^+ + B_s^- \longrightarrow FeB$$

The method makes use of the different energy level and capture cross section of FeB pairs and Fe interstitials. It is based on making two measurements of minority carrier lifetime: one in which the Fe atom is in a Coulombic pair and one in which it is not. Strong illumination is used in between measurements to easily dissociate FeB pairs [132]. For the lifetime measurement before the dissociation of FeB pairs, the effective lifetime, $\tau_{eff_{assoc}}$, can be written as:

$$\frac{1}{\tau_{eff_{assoc}}} = \frac{1}{\tau_{FeB}} + \frac{1}{\tau_{other\,assoc}} \tag{3.1}$$

where τ_{FeB} is the lifetime attributable to FeB pairs and τ_{other} takes into account recombination from all other mechanisms like surface recombination or boron-oxygen complexes. For the lifetime measurement after dissociation, assuming all FeB pairs are converted to interstitial iron, the effective lifetime, $\tau_{eff_{dissoc}}$, can be written as,

$$\frac{1}{\tau_{eff_{dissoc}}} = \frac{1}{\tau_{Fe_i}} + \frac{1}{\tau_{other\,dissoc}} \tag{3.2}$$

where τ_{Fe_i} is the lifetime attributable to recombination at Fe_i. If we can make the approximation $\tau_{other\,dissoc} \approx \tau_{other\,assoc}$, then the only difference in recombination between the two measurements is the change of impurity defect state from paired to

unpaired. By taking the difference in measured lifetimes in the dissociated and associated states and rearranging, the density of metal point defects can be determined by:

$$[Fe] = C\left(\frac{1}{\tau_{eff_{dissoc}}} - \frac{1}{\tau_{eff_{assoc}}}\right) \tag{3.3}$$

with the other terms in the effective lifetime canceling out, leaving only the Shockley-Read-Hall lifetimes for FeB and Fe_i which are proportional to the [Fe] concentration. The valuable insight of Zoth and Bergholz was that the above calculation for Fe point defect density is valid *even if other recombination mechanisms are active*, as long as the method used to break or form the FeB pairs does not change the recombination parameters of the other defects. For example, if the associated and dissociated measurements are taken at different injection levels (and not under low-injection conditions), then an Auger recombination term must be added to Eq. 3.3 to parameterize the known change in Auger recombination [133].

The pre-factor C can be determined based on the energy levels and capture cross sections of FeB and Fe_i for arbitrary injection conditions [134]. Since re-pairing of FeB beings immediately after dissociation, the fraction of iron that is in interstitial form, f, should also be taken into account for both the largely-associated and the largely dissociated cases. The C factor is then calculated by [135]:

$$C(\Delta n) = \frac{1}{f_{assoc}\cdot\chi^{FeB}_{assoc} + (1-f_{assoc})\cdot\chi^{Fe_i}_{assoc} - f_{dissoc}\cdot\chi^{Fe_i}_{dissoc} - (1-f_{dissoc})\cdot\chi^{FeB}_{dissoc}} \tag{3.4}$$

where:

$$\chi^{Fe_i} = \frac{\nu_{th}(N_A + \Delta n_{Fe_i})}{(1/\sigma_n^{Fe_i})(N_A + p_1^{Fe_i} + \Delta n_{Fe_i}) + (\Delta n_{Fe_i}/\sigma_p^{Fe_i})} \tag{3.5}$$

$$\chi^{FeB} = \frac{\nu_{th}(N_A + \Delta n_{FeB})}{(1/\sigma_n^{FeB})(N_A + \Delta n_{FeB}) + (1/\sigma_p^{FeB})(n_1^{FeB} + \Delta n_{FeB})} \tag{3.6}$$

where ν_{th} is the carrier thermal velocity, Δn is the injection level of the measurement, σ_p and σ_n are the capture cross sections of the defect for holes and electrons, respectively, and $p_1^{Fe_i}$ and n_1^{FeB} are the populations of holes and electrons when the Fermi level is at the defect energy level introduced by Fe_i and FeB, respectively, for

the calculation of SRH statisics (as in Eq. 2.1).

To be assured that the change in lifetime after dissociation is attributable to iron, the lifetime can be monitored as a function of time after dissociation of FeB pairs, with the lifetime evolving as the pairs reform. FeB repairing kinetics have been thoroughly investigated [136], with a time constant of reassociation given by:

$$\tau_{assoc} = \frac{5.5 \times 10^5}{N_A} T \exp\left(\frac{0.66\text{eV}}{\text{k}_B T}\right) \tag{3.7}$$

The activation energy of association, 0.68 eV, coincides with the activation energy of Fe diffusion in silicon [14], a probable indication that the FeB repairing reaction is wholly limited by Fe_i diffusion.

Table 3.1: Recombination Properties of Iron Point Defects for SRH Statistics

Defect	E_T (eV)	σ_n (cm^2)	σ_p (cm^2)	Reference
Fe_i	E_V + 0.38	4×10^{-14}	7×10^{-17}	[16]
FeB	E_C - 0.26	5×10^{-15}	3×10^{-15}	[137]

The analysis hinges on an accurate evaluation of the energy level and capture cross section of Fe_i and FeB. These values are updated with some regularity in the literature because of the importance of iron in silicon as a defect. The values used in this thesis are shown in Table 3.1.

3.2.2 Lifetime Measurement for Determination of Fe_i

To measure the concentration of iron in the point defect state, lifetime measurements are thus required in both the associated, FeB state and dissociation Fe_i state. Quasi-steady state photoconductance (QSSPC) was the primary method used to measure the minority carrier lifetime in this thesis [138]. QSSPC (WCT-120, Sinton Instruments) uses a flash lamp to illuminate the sample from above, injecting minority carriers. In general, the continuity equation for the average concentration of minority carriers in

the sample can be written as:

$$\frac{dn}{dt} = G(t) - \frac{\Delta n(t)}{\tau_{eff}(\Delta n)} \tag{3.8}$$

where $G(t)$ is the generation of carriers as a function of time, Δn is the concentration of injected carriers, and $\tau_{eff}(\Delta n)$ is the injection-dependent effective lifetime. In QSSPC, the flash illumination decays with a time constant (2.3 ms) much longer than the effective lifetime of most solar cell materials. Then, using a quasi-steady assumption, the generation can be considered to vary slowly enough that the system equilibrates, leaving simply:

$$\tau_{eff}(\Delta n) = \frac{\Delta n(t)}{G(t)} \tag{3.9}$$

The injection level $n(t)$ is calculated by measuring the change in conductivity of the sample using an inductively-coupled RF bridge beneath the sample. The injection level is related to the measured change in conductivity, $\Delta\sigma$ by:

$$\Delta\sigma = \mathrm{q}\Delta n(\mu_e + \mu_h)W \tag{3.10}$$

where q is the fundamental charge, μ_e and μ_h is the mobility of electrons and holes in the material respectively which are iteratively calculated from a mobility model, and W is the width of the sample. The inductive coil is sensitive over an area $\approx$2.5 cm in diameter, and the change in conductivity measured is an average over this area. The last unknown, the generation, is measured as a function of time using a calibrated reference Si cell. Plugging in the measured change in conductivity, all parameters are known in Eq. 3.9 and the effect lifetime can be extracted. The WCT-120 is also capable of operating in transient mode for samples with long lifetimes, or in a generalized mode where the full continuity equation (Eq. 3.8) is used to calculate the injection-dependent lifetime.

Interstitial iron measurements begin with the preparation of a known defect state by annealing on a hot plate for 10 min. at 200°C to avoid effects from the B-O defect [108]. The samples are then allowed to cool to room temperature. To dissociate

FeB pairs, the samples are exposed to 10 flashes of a 10 sun lamp [134]. Lifetime measurements were taken within 30 s after flashing. Measurements were then taken every 30 s for the next 3 min as FeB pairs begin to reform. Extrapolating the decaying lifetime back to that at t=0 after flashing allows for the calculation of the fully dissociated lifetime limited by Fe_i. The lifetime is then measured again after several hours depending on the resistivity of the sample to allow for a measurement in the fully re-associated FeB state. The interstitial iron concentration is then calculated using Eq. 3.3 and 3.4.

The same principles apply to the measurement of lifetime and interstitial iron by microwave photoconductance decay (μ-PCD), used in Chapter 6 to measure a spatially-resolved lifetime. In the μ-PCD measurement, a 200 ns laser pulse excites carriers in a 250 μm spot and the decaying lifetime transient (G=0) is measured by microwave reflectance. The laser is then rastered across the sample, mapping the effective lifetime. To achieve improved signal-to-noise, the measurement operates generally done under high, though unspecified, injection, leaving quantification of the C pre-factor for calculating interstitial iron concentration difficult if not impossible to accomplish. The spatially-resolved lifetime is valuable though, for analyzing heterogeneities in lifetime across samples.

The last technique used to measure lifetime and interstitial iron in Chapter 8 of this thesis is calibrated photoluminescence (PL) imaging [133]. In photoluminescence imaging of p-type materials, the intensity of the band-to-band radiation is related to the injection level by [139]:

$$I_{PL}(\hat{x}) \propto \Delta n(\hat{x})(N_A + \Delta n(\hat{x})) \tag{3.11}$$

where $\hat{x}$ denotes a particular local position on the sample. To calibrate PL intensity to absolute lifetime, a PLI measurement is first made under a known generation condition. Then, a QSSPC measurement is made, and the injection level in the sample region over the sensor is noted for the same generation as occurred in the PLI measurement. When done over a range of generation values (at least two),

the proportionality constants in Eq. 3.11 can then be extracted by averaging the PL intensity from the area that was over the sensor and comparing to the injection levels. The absolute lifetime can then be calculated from the PL intensities across the sample by $\tau(\hat{x}) = \Delta n/G(\hat{x})$. Finally, the wafers can be prepared for associated and dissociated lifetime measurements to calculate iron interstitial concentration as above. The advantage of PL is that it has high-spatial resolution and the injection condition can be well-calibrated.

3.3 Quantification of Iron Precipitate Distribution in Silicon by Synchrotron-based micro-X-ray Fluorescence

To build and evaluate impurity modeling of the distribution of iron after processing, we require accurate measurements of the initial, or as-grown, total iron concentration *and its spatial distribution*, as well as measurements of the iron distribution after processing to confirm the model. While readily available mass spectrometry methods can evaluate total bulk iron concentrations down to detection limits of single ppb [17], the heterogeneous, nanoscale spatial distribution of iron precipitates in silicon is more difficult to measure and quantify. Many traditional semiconductor characterization techniques, such as electron microscope-based electron beam induced current (EBIC), are limited in spatial resolution to a few microns [140]. Transmission electron microscopy (TEM) is capable of detecting isolated nanoscale particles [22] or agglomerations of precipitates [141], but due to the extensive (and destructive) sample preparation and small sampling volume, samples are generally intentionally contaminated for TEM analysis studies. Recently, a hierarchical electron-microscopy approach has been shown where EBIC is used to determine electrically-active defects, followed by focused-ion beam (FIB) to mill out the defective area and thin the lamella down, before TEM analysis of the structural and chemical properties [142]. Still, a representative (even statistical) sampling of the defects across the cm-scale device

remains to be desired.

Synchrotron-based x-ray microscopy techniques provide another characterization approach, previously used successfully to determine the elemental distributions of metal impurities and their chemical state within silicon solar cell materials [143, 144]. It has been shown that the effect of high-temperature processing involved in solar cell manufacturing on the redistribution of metals can be captured using synchrotron-based μ-X-ray fluorescence (μ-XRF) [145]. Leading X-ray fluorescence probes provide high photon fluxes and spot sizes on the order 100 nm or less, allowing for detection limits of approximately 10^{14} atoms/cm^2 per second of dwell time during mapping [41]. Fig. 3-1 illustrates the beam profile given by linescans on a 10 nm Cr film

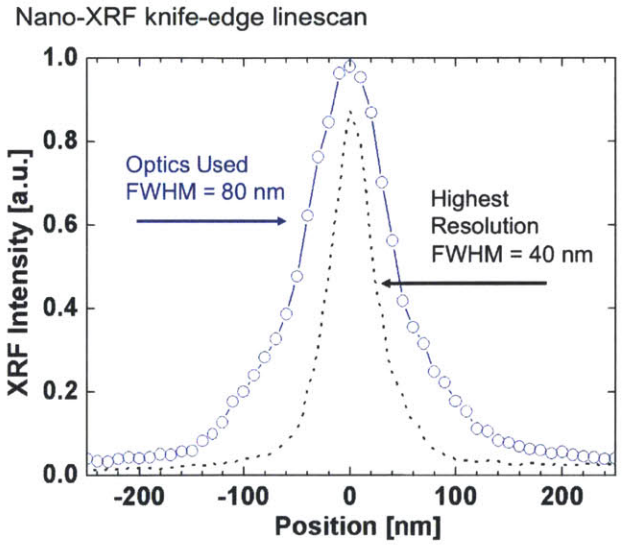

Figure 3-1: A scan of a Cr knife edge shows the full-width half maximum of the Hard X-ray nanoprobe at 26-ID of the Advanced Photon Source, using two different zone plate lenses. Reproduced from [146]

sectioned and scanned edge on at the Hard X-ray nanoprobe (HXN) operated by the Centre for Nanoscale Materials at Argonne National Laboratorys Advanced Photon Source beamline 26-ID (reproduced from [146]). The blue line corresponds to the zone plate used in this study (FWHM 80 nm). The dashed line demonstrates the highest resolution possible (FWHM 40 nm) with the hard X-ray nanoprobe using a state-of-the-art zone plate lens. To provide a higher photon flux, for the nano-XRF investigation shown in Chapter 8 we used a zone plate with lower resolution, higher

focusing efficiency, and similar lineshape for the trace element analysis. The efficiency for this zone plate relative to a 1000 mm^2 parallel beam at 8 keV is 4.34% with a focused flux of 5.2×10^8 photons/s.

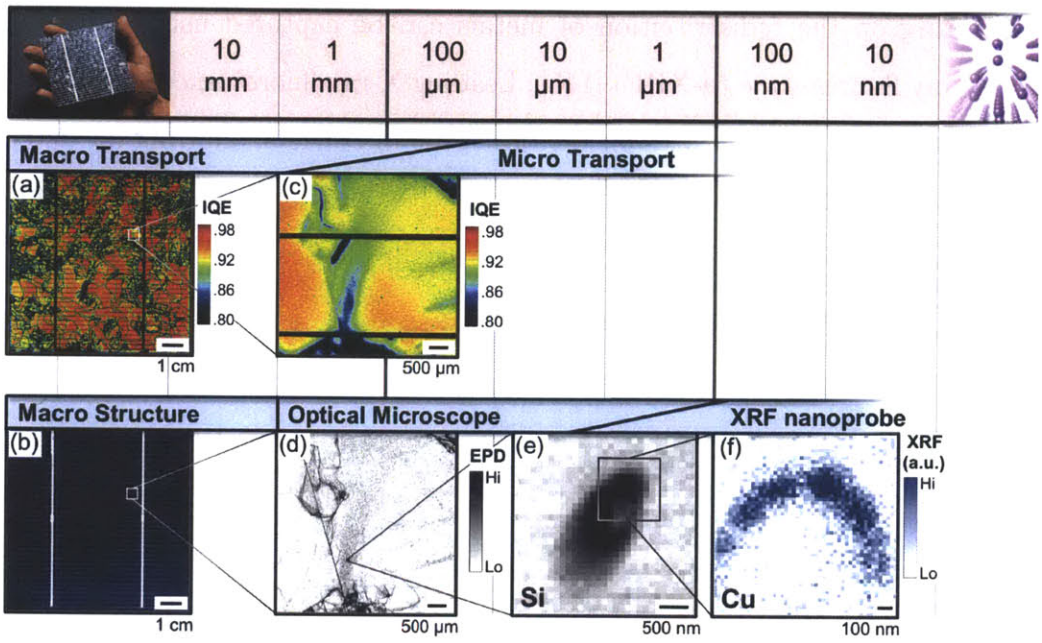

Figure 3-2: Hierarchical material characterization from centimeters to nanometers. (a) Laser-beam induced current and reflectivity mapping are combined to produce a full-wafer internal quantum efficiency (IQE) map, illustrating high- and low-performance areas. (b) The defects that reduce electrical performance are not visible in an optical image of the full cell. (c) High-resolution LBIC reveals neighboring high- and low-performance areas with micron-scale resolution. (d) Optical microscopy of the sample after defect-etching shows a high density of dislocations at highly recombination-active intragranular regions. (e) Synchrotron-based XRF is used to map the dislocation etch pits. The silicon channel is surface-sensitive and yields topographical information of etch pit location. (f) A region of the etch pit is selected for high-resolution nano-XRF, indicating that copper is present at the dislocation. Reproduced from [146]

Fig. 3-2 shows the hierarchical characterization approach using synchrotron-based XRF to bridge the characterization gap between micron scale optical or SEM techniques and atomic scale TEM. Fig. 3-2 shows the use of the nanoprobe to determine the elemental natures of nanoscale defects at intragranular dislocations in a commer-

cial mc-Si material. The results distinguish the metal-dislocation interactions that underly differences between benign and deleterious dislocations in a solar cell device and are detailed in Chapter 8.

While the sub-100 nm capability of the HXN represents the state-of-the-art, much of the XRF investigation in this thesis was completed at beamline 2-ID-D at the Advanced Photon Source of Argonne National Laboratory. There, the beamline uses a water-cooled Si(111) double-crystal monochromator and zone plate optics with 100 nm outer zone to achieve a circular spot size down to 150 nm (200 nm typical) with a focused photon flux of 4×10^8 photons/sec [147]. The principle advantage of running at 2-ID-D is higher throughput, allowing larger area measurements for better statistics.

3.3.1 μ-XRF Data Collection and Analysis

Fig. 3-3 shows the principles of XRF operation. An incident X-ray beam excites core shell electrons to the continuum, and the relaxation of outer shell electrons generates fluorescent photons which are collected by a detector. XRF provides sensitive chemical composition, since energy transitions within the atom are unique to each element. The primary fluorescent signal in the studies here comes from the K_α transition of an L-shell electron to the K-shell of the *3d* metal impurities under investigation. Other transitions are of course possible (K_β, L_α, etc.). Synchrotron light sources produce exceptionally bright sources of X-rays that allow for nano-focusing optics to be used. The sample can then be mounted on a nanopositioning stage so that the sample is moved relative to the beam, producing a point-by-point XRF map.

For sensitivity to the *3d* transition metals, in general in this work a 10 keV incident beam is used – just above the K_α absorption energy of Zn. An exception is made in Chapter 4, where for maximum sensitivity to iron, all samples except the phosphorus-diffused sister at the bottom of both ingots were measured in a single run using a 7.14 keV incident X-ray beam, just above the Fe Kα absorption edge. The advantage is increased sensitivity to iron, but we lose information about any heavier impurities. The phosphorus-diffused sister at the bottom of both ingots was measured at a subsequent run, and due to experimental constraints a 10.0 keV incident X-ray

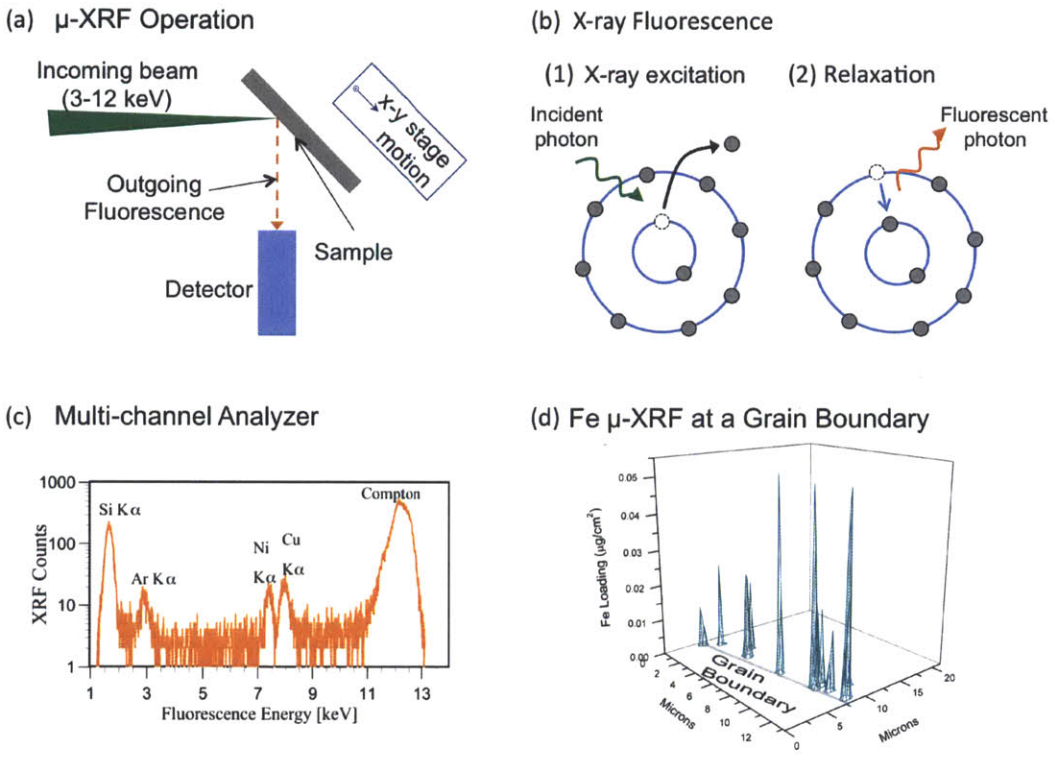

Figure 3-3: The principles of XRF operation are shown. (a) A sample is mounted on a positioning stage and scanned in front of an incident X-ray beam, producing fluorescent photons that are captured by a detector. (b) The incident X-ray energy is selected to exceed the binding energy of core shell electrons of the impurity of interest, exciting core shell electrons to the continuum (1), followed by the relaxation of outer shell electrons (2). To conserve energy, a fluorescent photon is produced during this relaxation transition. (c) The energy transition is elementally-specific, allowing sensitive identification of the impurity atom using a multichannel analyzer. (d) Synchrotron-based XRF thus allows the identification of small particles of impurities which are typically heterogeneously precipitated at structural defects, as is the case here for the iron particles seen decorating a grain boundary in multicrystalline silicon.

beam was used.

The beamline geometry is such that the sample is angled with respect to the incident beam with $\alpha = 15°$ and with respect to the detector such that $\beta = 75°$. The corresponding effective 1/e attenuation length of the fluorescent Fe signal in the μ-XRF setup was 6.6 μm for 7.14 keV incident radiation and 7.2 μm for 10 keV incident radiation, taking into account the angle of the incoming beam and the angle of the detector with respect to the sample normal. NIST standard reference materials 1832 and 1833 were measured in the same detector geometry in all cases to convert fluorescence counts into iron concentrations, using a peak fitting procedure by Vogt *et al* [148, 149]. The minimum detection limits for iron were 5 and 15.6 attograms for 7.14 and 10 keV, corresponding to a precipitate radius of 7.9 and 11.5 nm, respectively. Particle detection at the single-nanometer order of magnitude has been confirmed in separate experiments by the authors conducted at 2-ID-D where dots of metallic copper 8 nm in radius, produced by e-beam lithography, were detected on a silicon substrate with a 2:1 signal-to-noise ratio [150]. A typical map of Fe along a grain boundary in an as-grown multicrystalline silicon sample is shown in Fig. 3-4. This data will be used as an example for the remainder of the analysis.

3.3.2 Identification of Iron Precipitates

Following peak-fitting, analysis of the precipitate size distribution requires distinguishing the high-count Fe K_α signal due to precipitates from the low-level background noise across the map. The background noise can be mainly attributed to impurities in the detector and local photon scattering in the sample and off instruments in the beamline hutch. To isolate the iron-silicide precipitate data, the distribution of measured Fe counts for each map was fit with a truncated Gaussian distribution in order to estimate the distribution parameters of the background noise as shown in Fig. 3-5. A truncated normal distribution was assumed because of the physical impossibility of measuring iron counts in the detector system below 0. A maximum likelihood estimate was made for the mean, μ, and standard deviation, σ, of the measured Fe distribution, assuming the truncated normal distribution. A cutoff concentration was

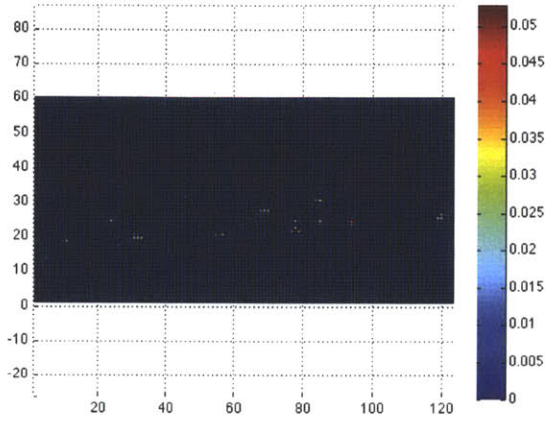

(a) Fe XRF Map

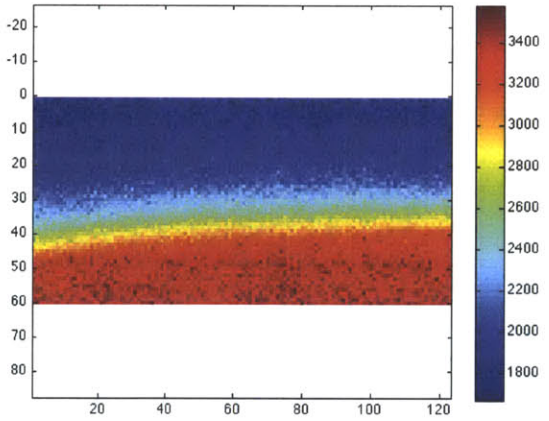

(b) Elastic Map

Figure 3-4: (a) 2D XRF map of Fe. Color (z-values) for Fe map in $\mu g/cm^2$. X and Y dimensions are in pixels. Pixels were 220 x 220 nm. (b) The corresponding elastic channel map is shown. Color in arbitrary units. High elastic contrast is seen across the grain boundary. The spatial distribution of Fe shows a strong segregation of Fe particles to the grain boundary region.

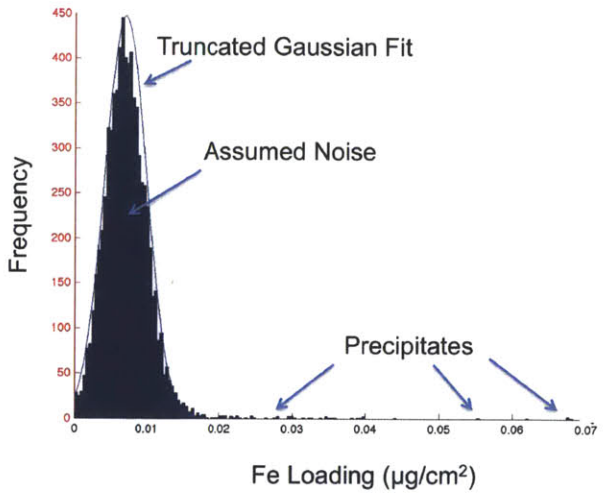

Figure 3-5: The distribution of Fe loading values from the XRF map in Fig. 3-4(a). The distribution is fit with a truncated normal distribution, and the mean and standard deviation are used to estimate a maximum value for the noise in the channel.

then taken to be the larger of the minimum detection limit or $\mu + n\,\sigma$, taking pixels with Fe counts below this as noise – pixels with iron concentrations above this cutoff were identified as containing precipitates. The number of standard deviations from the mean chosen for the cutoff, n, depended on the level of what could clearly be considered noise. In the data of Chapters 4 and 5, n was chosen to be 3.5. In Chapter 6 it was 3. Any pixel containing an Fe loading less than this noise cutoff was then considered noise, and I assume going forward that effectively all remaining high-Fe pixels above the cutoff are "signal."

3.3.3 Iron-Rich Particle Detection

A cluster of contiguous pixels above the noise cutoff was treated as a single precipitate when the highest iron concentration of any pixel within the cluster corresponded to a precipitate smaller than the step size (the case for all precipitates observed). The precipitate radius is then calculated directly from the number of iron atoms corresponding to that highest concentration pixel. This approach assumes that the high iron counts in surrounding pixels are attributable to bleed in from that central, large precipitate and avoids counting that large precipitate as many precipitates. To au-

tomate the process, individual particles are detected from the Fe fluorescence data of each map using a MATLAB image processing algorithm, *bwconncomp*, for binary images. The method takes the noise subtracted fluorescence data and finds components (i.e., precipitates) that contain non-zero pixels that are connected to at least 4 non-zero neighboring pixels. This avoids counting high-Fe pixels that bleed into neighboring pixels as being multiple particles. The particles detected using *bwconncomp* from the Fe Map in Fig. 3-4(a) are shown in Fig. 3-6.

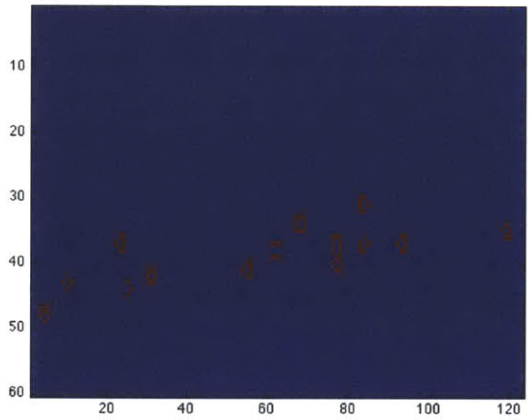

Figure 3-6: The particles detected from the Fe map in Fig. 3-4(a) are shown in the binary image above in red.

3.3.4 Calculation of Iron Precipitate Size and Density

To estimate the precipitate size, I first convert the scaled Fe loading value from $\mu g/cm^2$ at the surface to an effective number of fluorescing Fe atoms, N_{atoms}, by:

$$N_{atoms} = L \cdot \left(\frac{\pi * D^2}{4} \right) \cdot \frac{N_A}{MM_{Fe}} \tag{3.12}$$

where L is the loading of Fe in $\mu g/cm^2$, D is the full-width half maximum (FWHM) diameter of the incident beam in cm, N_A is Avogadro's number, and MM_{Fe} is the molar mass of Fe in $\mu g/mol$.

For particles calculated to be above the detection limit of the beamline, the number

of atoms per precipitate is then converted into an effective precipitate radius, R_{eff}, assuming the atoms are contained within a spherical $FeSi_2$ particle, according to:

$$R_{eff} = \left(\frac{3}{4\pi} \cdot N_{atoms} \cdot V \right)^{\frac{1}{3}} \tag{3.13}$$

where V is the volume of the primitive lattice of one Fe atom in $FeSi_2$, taken from the unit cell parameters in [151] .

The linear density of precipitates, Δ_{1D}, along the length of a grain boundary across a map is calculated by:

$$\Delta_{1D} = \frac{N_{pcps}}{GB\ Length} \tag{3.14}$$

where the GB length is also automatically calculated. Using the elastically-scattered data for the fluorescence map, shown in Fig. 3-4(b), I locate where the grain boundary exits the surface using image processing edge detection methods. Scattering intensities are observed to vary across a grain boundary because the crystal orientation determines the angular dependence of scattering (as in X-Ray Diffraction), and the detector subtends only a small solid angle. I detect this marked change in scattering intensity across the grain boundary within a map using a two-dimensional, Sobel discrete differentiation operator for edge detection in the measured scattering data. The application of the Sobel operator to the Elastic Map in Fig. 3-4(b) results in the binary image shown in Fig. 3-7(a). From the first pixels detected as edges according to the Sobel method in each line, I establish an estimate of the location where the grain boundary exits the surface of the sample by fitting the demarcated edge pixels with either a first-order or third-order curve, dependent upon whether the edge shows any curvature over the width of the map. The grain boundary curve determined by this fitting procedure is shown in Fig. 3-7(b). Line integration of the curve produces the estimate for *GBlength* needed in Eq. 3.14.

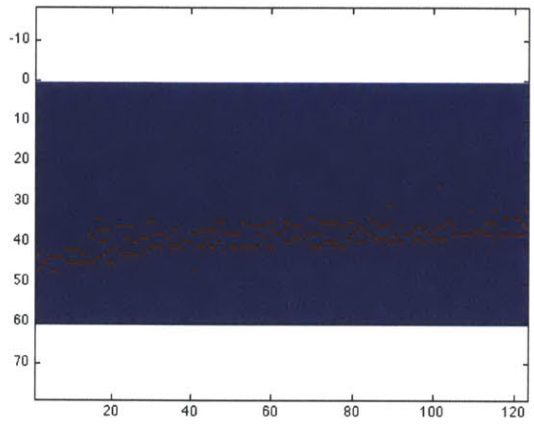

(a) Edge pixels in Elastic Map

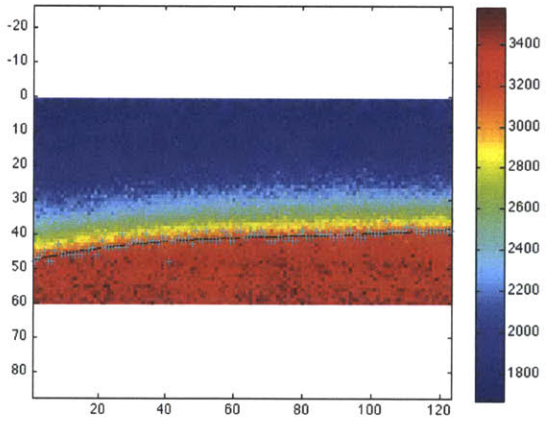

(b) Grain Boundary Fit

Figure 3-7: (a) The edge pixels in the Elastic Map from Fig. 3-4 as detected by the 2D Sobel operator. (b) Using the first edge pixel in each line (shown as crosses), a cubic function (in black) was fit to estimate the grain boundary location.

3.3.5 Assessment of Assumptions for Comparisons of XRF Data

Precipitation Dependence on Grain Boundary Type

In Chapter 4, we measure the distribution of individual precipitates along a small number of grain boundaries of several types, effectively conditioning the precipitate size distributions we measure at each GB on the GB type. The GB character potentially affects the measured precipitate distributions. In a study of the dependence of metal contamination on GB character, Buonassisi *et al.*[152] found that contamination levels measured by μ-XRF increased moderately with increasing GB energy, as characterized by coincident site lattice (CSL). Contamination levels were assessed by integrating the total metal XRF counts along a line scan across each GB. A high variance in decoration as a function of GB type was observed, especially at twin boundaries, where the variation in metal decoration was attributed to local dislocation density along the boundary, and at random angle grain boundaries – the two boundary types most prominent in this study. Similarly, Chen and Sekiguchi[153] used electron beam induced current (EBIC) to evaluate the recombination activity at GBs of different type in mc-Si samples intentionally contaminated on the wafer level. In good agreement with the μ-XRF of Buonassisi *et al.*,[152] they found that higher order CSL GB character and higher metal contamination lead to higher EBIC contrast. Bertoni *et al.*[154] found higher GB recombination activity after hydrogenation strongly depended on the dislocation density and/or faceting along the GB.

Several other authors have attempted to determine the impact of gettering on iron as a function of GB type. Arafune *et al.*[155] used μ-XRF with 2 μm resolution to take preliminary measurements of iron before and after gettering on a sample intentionally contaminated at the wafer level. Numerous, many-micron sized particles were observed relatively homogeneously distributed, much unlike the nanoparticles precipitated only along the GBs measured in this contribution. Tentatively, it was argued that small angle GBs trap more iron after gettering than the Σ3 GB investigated in that initial study. More recently, Takahashi *et al.*[156] used photoluminescence

contrast before and after gettering to evaluate the impact GB type has on gettering efficacy. They found that contrast between $\Sigma 3$ grain boundaries and the bulk grain generally decreased during gettering, suggesting that $\Sigma 3$ GBs do not inhibit the gettering process, while the photoluminescence contrast around RA GBs generally increased after gettering. However, comparisons of PL or EBIC contrast along a GB with respect to the bulk value before and after gettering are difficult to assess because bulk, intra-grain recombination often decreases substantially after a gettering step, meaning that the baseline for evaluating GB contrast changes dramatically between the two measurements before and after gettering. While this conditioning could shift the expected precipitate sizes of the different GB types, due to the small number of grain boundaries sampled and the generally high variance in the GB dependence data mentioned above, it is difficult to attribute changes in the precipitate size distributions here to changes in GB character.

Comparing adjacent wafers

Comparisons between as-grown and processed sister wafers, or wafers immediately adjacent to each other in the direction of crystal growth, must also be considered with some caution, because not only is the process condition changing, but the as-grown iron distribution can change as well, as seen in Chapter 4. Previously, it has been reported that sister wafers can have order of magnitude variations in the number of precipitates decorating the same GB passing through all samples.[157] However, this observation was made on sister wafers selected from the very top of an intentionally Fe and Cu contaminated ingot where the variation in the total metal concentration is strongest. In Chapter 6, we investigate wafers from less than 90% ingot height, and such large sister-to-sister differences are not seen between any two sisters regardless of ingot position, contamination level, or process condition. Furthermore, in this study we generally measure a larger area along each GB than previously,[157] which could help explain the smaller variation here.

Assumption of Precipitates at the Surface

Another important assumption in this analysis is that all precipitates are located at the surface of the sample. In reality, most precipitates lie along the plane of the grain bounday (GB) as it descends into the sample, leaving them detectable only because of the relatively large information depth of the hard X-ray beam. This finite particle depth has the effect of attenuating the effective Fe counts; thus the precipitate size estimates should be viewed as lower bounds. The depth effect should be relatively uniform across different grain boundaries due to the directional solidification of the ingot, which occurs largely perpendicular to the wafer surface.

3.4 Chapter summary

Of critical relevance to our work in iron gettering in silicon is the determination of the size and density distribution of iron precipitates in our experimental samples. μ-XRF and iron-boron pair dissociation have enabled unprecedented analysis of device-limiting defects in silicon solar cells. By applying iron-boron pair measurements and synchrotron-based X-ray fluorescence conducted at nanoprobes at Argonne National Laboratory's Advanced Photon Source, the large-sample, the quantitative measurement of the redistribution of metals during the high-temperature processing inherent in silicon solar cell manufacturing is enabled.

Chapter 4

Precipitated Iron As A Limit on Gettering Efficacy: Baselining a Standard Phosphorus Diffusion

The impurities-to-efficiency phosphorus diffusion gettering model is used to examine the efficacy of a standard gettering process on interstitial and precipitated iron in multicrystalline silicon. The model predicts a large concentration of precipitated iron remaining after standard gettering for most as-grown iron distributions. Although changes in the *precipitated* iron distribution are predicted to be small, the simulated post-processing *interstitial* iron concentration is predicted to depend strongly on the as-grown distribution of precipitates, indicating that precipitates must be considered as internal sources of contamination during processing.

To inform and validate the model, the iron distributions before and after a standard phosphorus diffusion step are studied in samples from the bottom, middle, and top of an intentionally Fe-contaminated laboratory ingot. A census of iron-silicide precipitates taken by synchrotron-based X-ray fluorescence microscopy confirms the presence of a high density of iron-silicide precipitates both before and after phosphorus diffusion. A comparable precipitated iron distribution was measured in a sister wafer after hydrogenation during a firing step. The similar distributions of precipitated iron seen after each step in the solar cell process confirm that the effect of standard gettering on precipitated iron is strongly limited as predicted by simulation.

Good agreement between the experimental and simulated data supports the hypothesis that gettering kinetics is governed by not only the total iron concentration but also by the distribution of precipitated iron. Finally, future directions based on the modeling are suggested for the improvement of effective minority carrier lifetime in multicrystalline silicon solar cells.

4.1 Introduction

Silicon solar cell performance outcomes remain strongly coupled to the execution of defect-engineering techniques during cell processing [75, 110]. Significant redistribution of iron and other metal impurities in silicon is possible during the high-temperature processing of silicon for solar cells – during crystal growth [17, 40], by phosphorus diffusion [70, 97], by extended anneals at temperatures where defects

remain mobile (*i.e.*, low-temperature annealing) [33–37, 60], and even during metallization firing [158].

In this chapter, a cohesive assessment is presented of the redistribution of total iron concentrations – both precipitated and interstitial – using modern industrial processing steps and relate it to the final impact on device performance. We simulate the coupled behavior of phosphorus and iron during cell processing using the Impurities to Efficiency (I2E) model [103] and support the model with a nanoscale experimental investigation of the evolution of iron during solar cell processing using synchrotron-based micro-X-ray fluorescence (μ-XRF). To test the impact of the firing step on the distribution of precipitated metals, we also examine the effect of a faux-firing step on gettered samples, where the wafer is annealed according to the time-temperature profile of a metallization firing but without the presence of external metals for contacting, to isolate the impact of internal metal redistribution. We explain our empirical results using the model, identifying the gettering limitations of standard industrial processing. Finally, we assess the broader gettering parameter space to identify paths toward iron-tolerant silicon solar cell processing.

4.2 Simulation of a Standard Phosphorus Diffusion Gettering Process

First, simulations were performed to study the impact of a standard phosphorus diffusion step on the concentration of interstitial iron, Fe_i, and on the concentration of precipitated iron, Fe_p, in form of β-iron-silicide, β-$FeSi_2$. The sum of both concentrations is referred to as *total* iron concentration, which is approximately equal to the Fe_p concentration because the latter is usually about two orders of magnitude higher than the Fe_i concentration [123].

In this work, we simulated a standard phosphorus diffusion of 15 min at 850°C followed by a 10 min free cooling with an exponential time constant of 7 min. These conditions should approximate the phosphorus diffusion profile used to process the

wafers used in this study, which originate from the European Integrated Project CrystalClear as described in Section 4.3. With this time-temperature profile, simulations were performed for a typical range of total as-grown iron concentrations in silicon solar cells from 2×10^{13} to 1×10^{15} cm^{-3}, and a typical range of as-grown β-$FeSi_2$ precipitate radii from 8 to 18 nm [42]. The as-grown Fe_i concentration, which has little impact on final Fe_i concentration when a large fraction of total iron is precipitated, was assumed to be 1×10^{13} cm^{-3} for all simulations.

4.2.1 Factors determining interstitial iron concentration after gettering

Iron point defects are thought to be the most recombination active form of iron, given their diffuse distribution throughout the material and mid-gap energy level [16].

Fig. 4-1(a) shows the post-processed Fe_i concentration as a function of these two main iron distribution dependencies: the total as-grown Fe concentration (approximately equal to the as-grown precipitated iron concentration) and the as-grown average precipitate radius. The effect of the precipitated iron distribution on dissolution comes from the dependence in Eqn. 2.9 on the density of precipitates, N_p, and precipitate radius, r. By substituting N_p for its fundamental constituents, the as-grown precipitated iron concentration, C_p, the volume of one iron atom in $FeSi_2$, $V_{Fe/FeSi_2}$, and the volume of a single spherical precipitate of average as-grown radius, r_o, we have more explicitly,

$$\frac{dC_i}{dt} = 4\pi \frac{C_p \cdot V_{Fe/FeSi_2}}{\frac{4}{3}\pi r_o^3} \cdot r D_{Fe} \left(C_{eq} - C_i\right). \tag{4.1}$$

This reduces to:

$$\frac{dC_i}{dt} \propto \frac{C_p}{r_o^2}\frac{r}{r_o} \cdot D_{Fe} \left(C_{eq} - C_i\right). \tag{4.2}$$

Across the entire parameter space, the post-processed Fe_i concentration varies between about 2×10^{8} and 1×10^{12} cm^{-3} and shows a strong dependence on both as-grown parameters. Three distinct regions with respect to as-grown total iron concentration

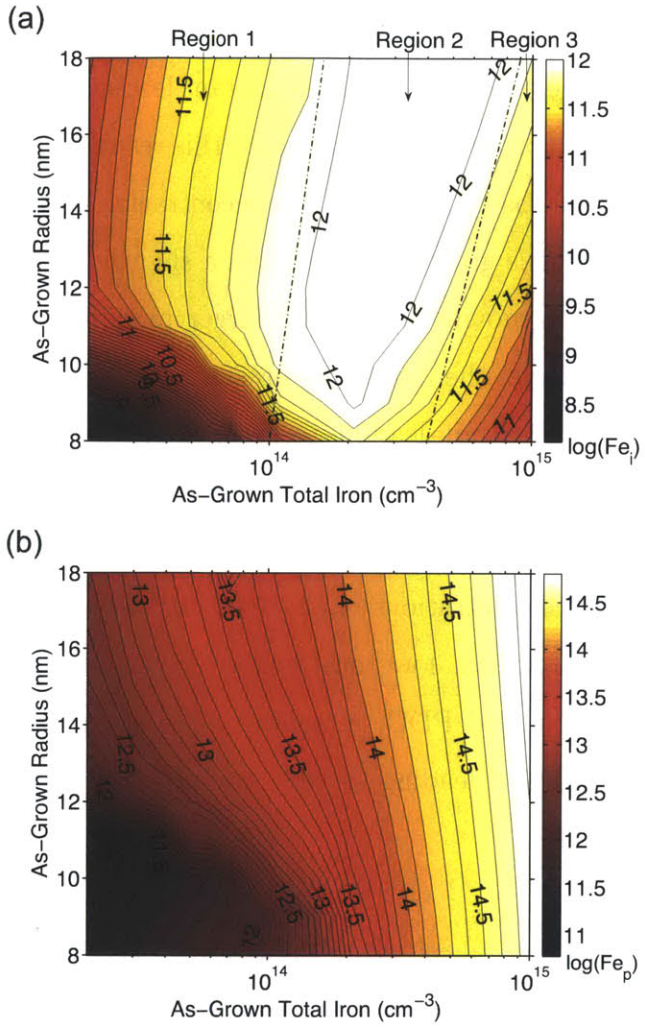

Figure 4-1: Simulation results are shown for a 15 min phosphorus diffusion at 850°C with 10 min cooldown: (a) Fe_i concentration and (b) precipitated Fe concentration in the wafer bulk after gettering as a function of the total as-grown Fe concentration and as-grown iron-silicide precipitate radius; please note the different scale bars. In (a), thin doted lines demarcate three different regions of Fe_i response to a standard phosphorus diffusion gettering process. Reproduced from [159].

are observed in Fig. 4-1(a) and are described in short below. Further discussion of these regions can be found elsewhere [11, 120].

Region 1:

Starting at as-grown total Fe concentrations $\lesssim 2\times10^{14}$ cm^{-3}, the post-processed Fe_i concentration decreases moving left in Fig. 4-1(a) toward smaller as-grown total con-

centrations. As can be seen from Eqn. 4.2, decreasing total Fe concentration while keeping as-grown average precipitate radius constant leads to a decelerated dissolution during high-temperature processing. Ultimately, this results in lower final Fe_i concentration. Similarly, in this as-grown total Fe concentration region of $\lesssim 2\times10^{14}$ cm^{-3}, when the as-grown radius is greater than 12 nm, a slightly decreasing post-processed Fe_i concentration is seen moving up toward larger precipitate sizes along as-grown isoconcentration lines. Here, the increasing precipitate radius causes decreased dissolution, as seen in Eqn. 4.2, again leading to lower final Fe_i concentration.

For smaller radii in Region 1, however, the post-processed Fe_i concentration decreases with decreasing precipitate radius, and minimum Fe_i concentration is found for total as-grown iron concentrations of $(2-3)\times10^{13}$ cm^{-3} and a precipitate radius of 8 nm. In this small radius, low total concentration regime, almost all precipitates dissolve during standard PDG, removing some of the kinetic limitation of Fe extraction from the bulk and leading to very small final Fe_i concentrations.

Region 2:

For higher as-grown total Fe concentrations between 1×10^{14} and 4×10^{14} cm^{-3}, the dissolution of precipitates during high-temperature processing partially offsets the Fe_i reduction due to segregation gettering, and a relatively high post-processed Fe_i concentration 1×10^{12} cm^{-3} is obtained after standard PDG.

Region 3:

For an as-grown Fe concentration $\gtrsim 4\times10^{14}$ cm^{-3}, a decrease of the post-processed Fe_i concentration is observed with increasing as-grown Fe concentration and with decreasing as-grown radius. In this region, the density of precipitates increases such that the linear distance between precipitates decreases down to a few microns. With decreasing distance between precipitates, an enhanced precipitation of Fe_i atoms takes place during cool down to room temperature, and a lower post-processed Fe_i concentration is obtained due to an internal gettering effect. While internal gettering plays an increasing role at very high precipitate densities, we likely overestimate this

internal gettering effect because we assume a homogeneous distribution of iron precipitates in our model. In reality, precipitates are heterogeneously distributed at favorable nucleation sites such as grain boundaries and dislocations [17, 146], such that intragranular precipitates sites may be much further separated than we assume here, leading to decreased internal gettering interactions.

4.2.2 Factors determining precipitated iron concentration after gettering

The vast majority of iron atoms in as-grown material are found in $FeSi_2$ precipitates. However, because of their low spatial density, $FeSi_2$ precipitates usually have a smaller direct impact on lifetime than Fe_i. Still, $FeSi_2$ precipitates indirectly affect lifetime by dissolving during high-temperature processing, introducing Fe_i into the bulk. In conventional solar cell processing, phosphorus diffusion is the step with the largest thermal budget and therefore the largest potential to control the distribution of iron.

The corresponding post-processed Fe_p concentration as a function of the total as-grown Fe concentration and the as-grown precipitate radius is shown in Fig. 4-1(b). It varies between about 1×10^{11} and 7×10^{14} cm^{-3} across all input values. For all precipitate radii, the post-processed Fe_p concentration decreases with decreasing total as-grown iron concentration. The largest decrease in the Fe_p concentration is observed for the smallest precipitates, and a minimum Fe_p concentration is found for total as-grown iron concentrations of $(2-3)\times10^{13}$ cm^{-3} and precipitate radii of 8–11 nm.

It is observed that only a small fraction of Fe_p is removed during gettering toward the highest total iron concentrations considered. This can be explained by two simple estimates: the solid solubility of Fe in Si at 850°C is only about 1.2×10^{13} cm^{-3} [43], *i.e.*, about two orders of magnitude lower than the highest total iron concentrations typically found toward the top and borders of mc-Si ingots [27, 123]. The Fe_i diffusivity in Si at 850°C is about 9.8×10^{-7} cm^2/s [14], *i.e.*, an Fe_i atom needs about 80 s to diffuse from the center of a 180 µm thick wafer to the gettering layer at the wafer

surface. In conclusion, at 850°C, both the solid solubility and the diffusivity strongly limit the effective reduction of precipitated iron in the wafer bulk.

4.3 Intentionally Fe-Contaminated Sample Specifications, Processing, and Measurement Details

To investigate these simulation results experimentally, samples were selected from the bottom, middle, and top of two laboratory-scale ingots intentionally contaminated with 53 and 200 ppmw Fe as part of the European Integrated Project CrystalClear. Extensive data on the macroscopic distribution of iron and its effect on performance in these samples has already been published [21, 26, 160]. The samples were selected from a relative ingot height of 29%, 61%, 88% and 30%, 68%, and 85% for the 53 and 200 ppmw Fe ingots, respectively. The wafers from these heights will subsequently referred to as "Bottom," "Middle," and "Top." The samples from the bottom were selected $\gtrsim$3 cm above the crucible bottom to avoid the regions where contamination from solid-state in-diffusion from the crucible dominates the metal distribution. Interpolating the neutron activation analysis (NAA) data from the CrystalClear project [26, 160], the iron concentrations for the samples can be estimated to be: $3.3\times10^{13}, 3\times10^{13}, 6\times10^{14}$ for the 53 ppmw ingot and $5\times10^{13}, 1.2\times10^{14}, 4.4\times10^{14}$ atoms/cm^3 for the 200 ppmw ingot.

Three sister wafers were pulled out of the processing line at each of the three heights in both ingots: one as-cut, one after gettering, and one following SiN_x deposition and firing [21] . The phosphorus diffusion process was a standard industrial process conducted in a $POCl_3$ tube furnace. For the hydrogenation and firing step, the same temperature profile was carried out as for the metallization firing, but no metals were present so as to avoid contamination effects. A schematic of the sample selection from the iron-contaminated ingots is shown in Fig. 4-2.

For each set of sister wafers from the six different locations, a single grain boundary

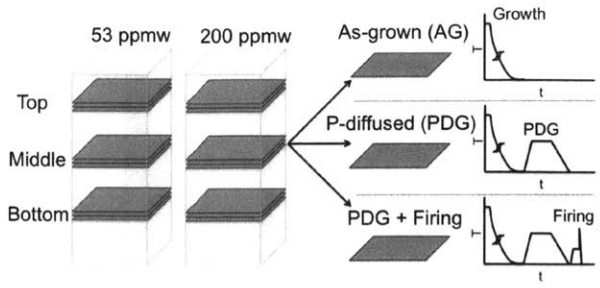

Figure 4-2: The samples for the μ-XRF study were selected from three heights in ingots of low and high iron contamination. The sister wafers were pulled out of the process after significant high-temperature steps for comparison. Reproduced from [159].

Table 4.1: Properties of the samples analyzed using synchrotron-based μ-XRF. Estimated bulk iron concentration is shown, as is the type of grain boundary analyzed by μ-XRF.

Ingot	Height (%)	[Fe] (cm^{-3})	GB type
53 ppmw Fe	29	3.3×10^{13}	RA (38.1°)
	61	3×10^{13}	Σ3
	88	6×10^{14}	RA (38.7°)
200 ppmw Fe	30	5×10^{13}	Σ27*b*
	68	1.2×10^{14}	RA (27.2°)
	85	4.4×10^{14}	Σ3

(GB) appearing in each sister was selected for μ-XRF.

The GB character and misorientation angles, analyzed by electron back-scatter diffraction measurements, are reported for the different sample groups in Table 4.1. The regions of interest were selected from the same 2.5 cm × 2.5 cm area within all wafers, but the grain structure varied moving between the widely separated ingot heights, resulting in differing GB type. High-order coincidence site lattice and random angle GB character have been associated previously with higher degrees of impurity decoration, although large variations exist within the data [152]. The potential impact of different GB type on the precipitation behavior is discussed in detail in Chapter 3.

Before μ-XRF measurement, all samples were cleaned in organic solvents, followed by a cleaning step with HCl. Beamline 2-ID-D at the Advanced Photon Source at

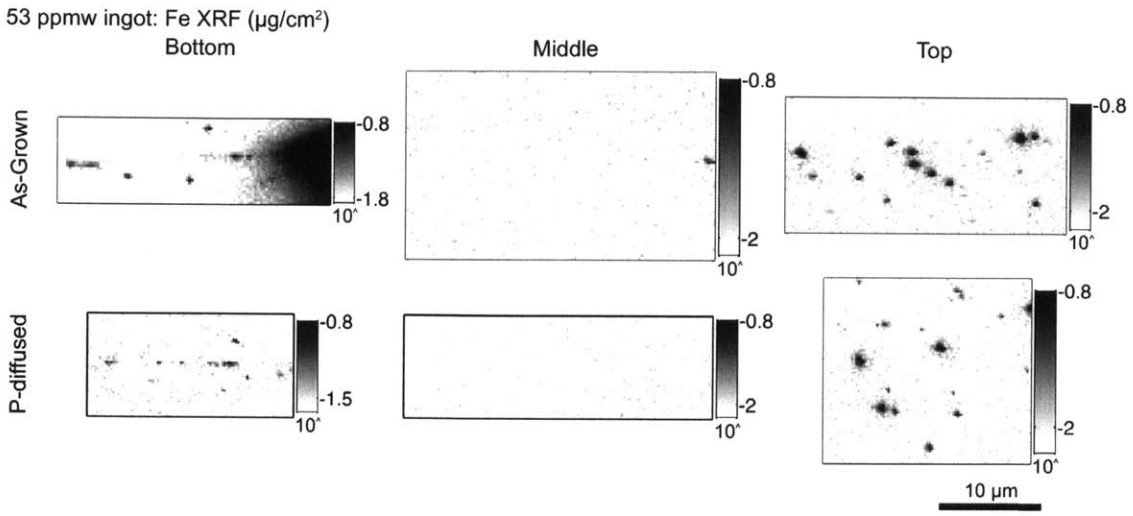

Figure 4-3: As-grown and phosphorus-diffused distribution of iron in the 53 ppmw ingot, as measured by μ-XRF at approximately the same grain boundary position on sister wafers at three ingot heights. Pixel size is 220 nm x 220 nm. Iron concentrations are shown in units of $\mu g/cm^2$. Concentrations are plotted in logarithmic scale to enhance the contrast between the dark iron-silicide precipitates and the lighter background level. The bottom and top of the ingot show significant decoration of grain boundaries, while the middle shows minimal decoration. Reproduced from [159].

Argonne National Laboratory was used to collect the μ-XRF measurements [147], utilizing its 200 nm full-width half-maximum spot size and taking 220 nm steps during μ-XRF mapping with a 1 second dwell time per pixel. An effective precipitate radius was calculated assuming all iron atoms within high-Fe pixels could be attributed to a single spherical β-$FeSi_2$ precipitate located at the surface of the sample. Further details on the data analysis and the extraction of precipitate distributions from the μ-XRF measurements can be found in Chapter 3.

4.4 μ-XRF Results on As-Grown, $POCl_3$ Gettered, and Hydrogen-Passivated Samples

The μ-XRF maps collected on as-grown and phosphorus-diffused sister samples from the bottom, middle, and top of the 53 ppmw ingot are shown in Fig. 4-3. Fig. 4-4 shows the corresponding maps for the 200 ppmw ingot. The μ-XRF maps were taken at approximately the same position (within 100 μm) along the same grain boundary

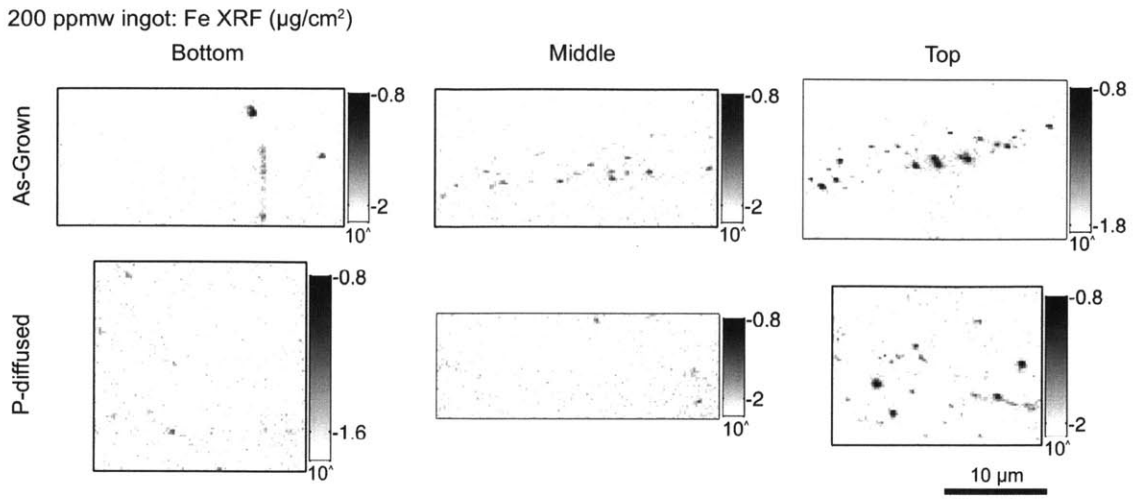

Figure 4-4: As-grown and phosphorus-diffused distribution of iron in the 200 ppmw ingot, as measured by μ-XRF at approximately the same grain boundary position on sister wafers at three ingot heights. The observed precipitate size and density appear to increase along the ingot height. Reproduced from [159].

in the sister wafers from a given height. The small changes in position along the grain boundary account for the differences in the size and shape of the μ-XRF maps between the as-grown and phosphorus-diffused sisters. Because the samples were mounted such that each GB of interest was horizontal, most maps show a largely horizontal line-up of precipitates along the GB.

In the 53 ppmw as-grown samples, significant precipitation is observed at the bottom and top of the ingot, while the GB in the middle of the ingot was relatively clean, exhibiting only a single large, readily observable precipitate along its length. A slightly higher density of precipitates is observed at the top of the ingot with respect to the bottom. After phosphorus diffusion, a number of precipitates remain along the GB for both the bottom and the top. In the middle of the ingot, no precipitates are immediately observable after phosphorus diffusion.

For the 200 ppmw as-grown samples, increasing GB decoration is seen moving from the bottom to the top of the ingot. Note that in the map of the as-grown sample taken from the bottom of the 200 ppmw ingot, a vertical step was observed in the GB, resulting in heavy precipitation there. In the phosphorus-diffused sisters, a few precipitates are seen at the bottom (though faint) and at the top of the ingot.

In the middle of the ingot, few precipitates are visible after gettering.

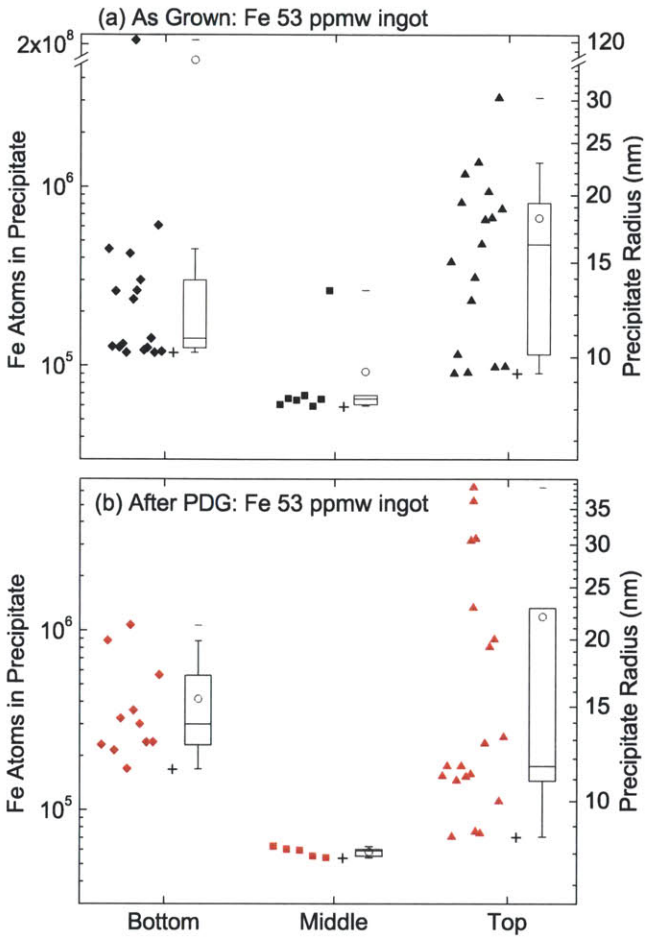

Figure 4-5: As-grown distribution of iron in the 53 ppmw ingot (a) before and (b) after gettering. Boxplots show the quartiles of the distribution, and the open circles the mean. The cross underneath the data shows the noise cutoff value from the automated analysis. The bottom and top of the ingot show significant decoration of grain boundaries, while the middle shows minimal decoration. The distribution after gettering is similar to the as-grown distribution. Reproduced from [159].

To achieve a more quantitative comparison of the distribution of precipitate sizes measured in each sample, iron-silicide precipitates were isolated within the maps and their concentrations converted to effective $FeSi_2$ radii values as detailed in Chapter 3. A quantification of precipitate size distribution is presented in Fig. 4-5 for the 53 ppmw ingot and Fig. 4-6 for the 200 ppmw ingot. Alongside the data, a boxplot of their distribution shows the 25%, 50%, and 75% percentiles. Open circles indicate the mean of the distribution for each sample. The cross at the bottom of each data set

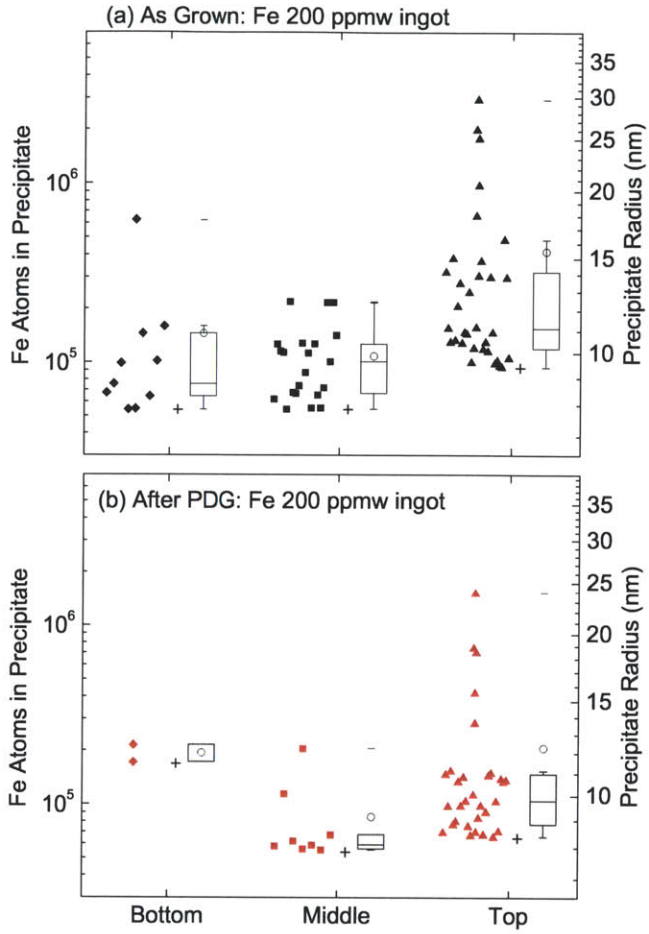

Figure 4-6: As-grown distribution of iron in the 200 ppmw ingot (a) before and (b) after gettering. The median as-grown precipitate size increases along the ingot height. The observed precipitate density also increases with ingot height. After phosphorus diffusion, the distributions are shifted toward smaller values at the middle and top of the ingot. Reproduced from [159].

indicates the automated noise cutoff value for each sample, details about which can be found in Chapter 3. Table 4.2 contains the measured linear density of precipitates along the GB in each region of interest.

For the as-grown 53 ppmw samples, the top shows the largest median precipitate size, as expected, corresponding to higher iron concentrations toward the top of the ingot measured by NAA by Kvande *et al* [26]. In the middle of the ingot, the automated analysis reveals a single large precipitate of around 3×10^5 iron atoms and suggests several precipitates of around 7×10^4 atoms in size. As these smallest precip-

Table 4.2: The linear density of precipitates per micron at each GB measured is shown. At the top of the 53 ppmw ingot, an additional sister sample was measured after a firing step – its linear precipitate density is shown in parentheses under the P-diffused column.

Ingot	Height (%)	As-Grown Pcp. Density (μm^{-1})	P-Diffused Pcp. Density (μm^{-1})
53 ppmw Fe	29	0.57	0.56
	61	0.28	0.16
	88	0.63	0.96 (1.08)
200 ppmw Fe	30	0.35	0.10
	68	0.78	0.30
	85	1.27	1.53

itates are near the noise cutoff value, are also close to the fundamental detection limit of our setup, and are not readily observable in Fig. 4-3, it is difficult to say definitively whether they are attributable to noise or represent an accurate measurement of precipitated iron. The bottom of the ingot reveals a generally narrower distribution of sizes compared to the more heavily contaminated top of the ingot.

The precipitate size distributions after phosphorus diffusion in the 53 ppmw ingot are very similar to those in the as-grown samples, as revealed by a direct comparison of Fig. 4-5(a) and 4-5(b). Sizes at the bottom again range from about 10^5 to 10^6 Fe atoms, while at the top there is a wider distribution from $7{\times}10^4$ to almost 10^7 Fe atoms. In the middle sample, no large precipitates were observed after phosphorus diffusion.

In the 200 ppmw ingot, the as-grown precipitate distribution shifts distinctly to larger values moving from the bottom to the top of the ingot as the metal contamination level increases, as seen in Fig. 4-6. After diffusion, the range of precipitate sizes shifts to smaller values particularly at the top of the ingot as can be seen comparing Fig. 4-6(a) to 4-6(b). Additionally, in the phosphorus-diffused middle sample, gettering appears to decrease the number of precipitates larger than $\gtrsim 7 \times 10^4$. After gettering, the bottom sample revealed only a small number of precipitates, in part due to the higher minimum detection limit in that sample, though the mean size was

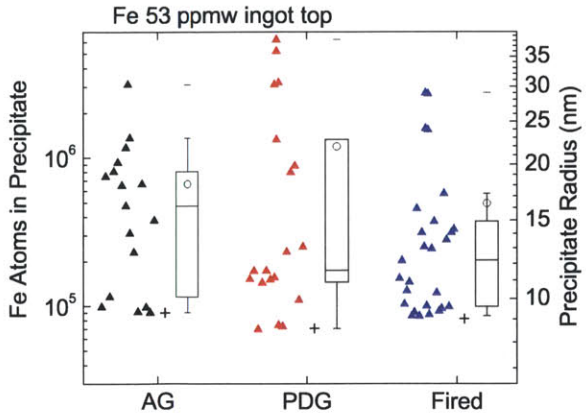

Figure 4-7: Direct comparison of the three sister wafers with different processing from the top of the 53 ppmw ingot: as-grown (AG), phosphorus-diffused (PDG), and fired. Significant amounts of iron remain precipitated at grain boundaries after P-diffusion, potentially acting as sources for point defects during later processing. The time-temperature profile of the hydrogenation step and subsequent faux-firing does little to change the precipitated iron distribution. Reproduced from [159].

similar to the as-grown value.

Although inferences from the precipitate density data should be taken with caution due to the small number of GBs investigated (one at each ingot position), there are several further notable observations from the 200 ppmw ingot. First, the overall density increases with increasing relative ingot height as seen in the data of Table 4.2. Finally, phosphorus diffusion does not decrease the density of precipitates at the top of the ingot, but does reduce precipitate density in the bottom and middle of the ingot.

To test the re-distribution of precipitated metals during anti-reflection coating deposition and subsequent firing steps (with no metals present), we analyzed the metal distribution at the GB region of interest on a third sister wafer from the top of the 53 ppmw ingot. While phosphorus diffusion gettering offers the most significant opportunity for impurity redistribution because of its high temperature and duration, subsequent firing for metallization has been shown to have an impact at least on interstitially-distributed metals despite its very short extent [158, 161]. The precipitate size distributions in the as-grown, phosphorus-diffused, and fired sister

wafers from the top of the 53 ppmw ingot are shown in Fig. 4-7. As can be seen quite clearly, a vast majority of the as-grown precipitated iron remains at the end of solar cell processing. The range of precipitate sizes and the density at the top of the ingot are largely unaffected by any of the thermal processing.

4.5 Discussion

4.5.1 As-Grown Iron-Silicide Precipitate Size Increases with Ingot Height

Iron-silicide precipitates must increase in size, density, or both along the ingot height to follow the increase in total iron concentration due to solid-liquid segregation during crystal growth. From the data of Figs. 4-5(a) and 4-6(a), it appears that the median precipitate size increases with increasing ingot height. In the 53 ppmw ingot, the as-grown median precipitate size in the top is larger than at the bottom (see Fig. 4-5(a)). The middle of the 53 ppmw ingot yielded such a small number of definitively identified precipitates that it makes it difficult to use for comparison. The increase in precipitate size along the ingot height is clearer in the 200 ppmw ingot (Fig. 4-6(a)).

We can compare the experimental as-grown size distributions to recent theoretical efforts to predict the heterogeneous nucleation and growth of iron-silicide precipitates during crystallization. Haarahiltunen *et al.* [40] simulated the average precipitate radius after cooling from solidification at 1°C/min, a cooling rate comparable to the cooling rate of about 2°C/min that was applied to the ingots of our investigation [26]. In general agreement with the experimental data herein, they found that the average precipitate radius increases for increasing total iron concentration. They calculated an average precipitate radius that varied between 5 and 180 nm for as-grown iron concentrations between 10^{14}atoms/cm^3 and 1×10^{15}atoms/cm^3. However, the precipitate radii experimentally observed here are generally much smaller than those simulated, ranging only between the 7.9 nm detection limit and roughly 30 nm, excluding an outlier in the bottom of the 53 ppmw ingot at roughly 120 nm radius.

Larger as-grown average precipitate radii may be found in larger commercial mc-Si ingots. While significant effort was made to ensure that the solidification environment of the laboratory-scale ingots used here was representative of industrial growth [26], the large difference in ingot mass (12 kg. here vs. 650 kg. typical for Gen6 furnace) resulted in a faster cooling rate in the small, laboratory-scale furnace – a cooling rate of slightly less than 2°C, as mentioned above. As shown by simulation [162], an increase in cooling rate leads to metal impurity supersaturation developing more quickly, leading to more widespread precipitation because of the larger supersaturation driving force. With nucleation and precipitation occurring at a higher density of sites, ultimately smaller precipitate sizes are expected, as has been observed experimentally for even faster cooling rates of sheet and ribbon silicon materials in comparison with ingot cast multicrystalline silicon [17].

Similarly, a larger supersaturation is assumed to develop in ingots with higher contamination levels when cooled with the same temperature profile, leading to more widespread nucleation and formation of smaller precipitates. Indeed we observe that the median precipitate sizes for the bottom and top of the 200 ppmw as-grown samples are smaller compared to the 53 ppmw as-grown samples.

4.5.2 As-Grown Iron-Silicide Precipitate Density Varies with Ingot Height

The precipitate density appears to increase along the ingot height as well, seen most clearly for the 200 ppmw ingot in Table 4.2. A similar trend was found by nucleation and growth simulation by Haarahiltunen *et al* [40]. Schön *et al.* [49] have also simulated the nucleation and growth of iron-silicide precipitates in wafers with an iron concentration of 3.5×10^{13} atoms/cm^3. In such wafers, they expect a much higher density of small precipitates ($10^3 - 5 \times 10^5$ atoms) than large precipitates ($> 5 \times 10^5$ atoms). With an experimental detection limit of $\gtrsim 5 \times 10^4$ atoms, we still measure a larger density of small precipitates ($< 5 \times 10^5$ atoms), though not the two orders of magnitude difference predicted in their study.

Conclusions from the experimental precipitate density data are tentative, however, as the number of as-grown grain boundaries measured to calculate density is small (6, from the different ingot locations investigated). Experimental precipitate size distribution data, on the other hand, is more robust because a large sample of precipitates were measured in the as-grown state (> 100). Additional details on the assumptions involved in the extraction of the precipitate distribution data can be found in Chapter 3. In short, future studies of the changes in the as-grown iron-silicide precipitate distribution as a function of iron concentration and cooling rate are needed.

4.5.3 Iron-Silicide Precipitate Distribution after Standard Processing

The typical phosphorus diffusion step applied to these samples failed to remove a significant fraction of the precipitated iron, particularly at the top of the two ingots. Fig. 4-7 summarizes this experimental finding. This lack of significant reduction of precipitated iron is to be expected from the simulations of Fig. 4-1, because at the top of both ingots the as-grown total iron concentrations and the measured precipitate size distribution position the samples near the border between Region 2 and Region 3. In this parameter space, standard gettering is predicted to be largely ineffective at decreasing precipitated iron. In accordance with this modeling result, at the top of the 53 ppmw ingot we observe that mean precipitate size is actually higher for the phosphorus-diffused sister when compared with the as-grown sister, though the distribution remains widely scattered, and a reduction in the median precipitate size is seen. An increase in mean precipitate size after gettering has been seen experimentally and in simulation previously for heavily contaminated material for the edge of an ingot [42, 162]. Likewise, at the top of the 200 ppmw ingot, the distribution is largely unchanged, with a small shift toward smaller precipitates.

The efficacy of the standard diffusion step at removing precipitated iron varies, however, depending on the iron distribution parameters. The most successful reduc-

tion in precipitated iron seems to occur in the middle of both ingots, where the median precipitate size and the precipitate density both decrease in the phosphorus-diffused sister, although statistics are poor, especially with respect to large, readily observable precipitates. With the poor statistics qualifying any strong conclusions, this reduction in the middle of the ingot does agree well with the simulations of Fig. 4-1(b). The as-grown iron distribution measurements in the middle of the 53 ppmw (8 nm median radius, 3×10^{13} atoms/cm^3 total Fe) put that sample well within the small-radius, low iron regime of Region 1 where precipitated iron dissolution and removal can be effective.

For the middle of the 200 ppmw ingot, the as-grown iron distribution measurements (10 nm median radius, 1.2×10^{13} atoms/cm^3 total Fe) position it within Region 2. Because of the small median radius, the simulations of Fig. 4-1(b) indicate that standard gettering begins to remove substantial amounts of precipitated iron. We observe reduction in the median iron-silicide precipitate size experimentally, as seen in Fig. 4-6.

The difference in the iron distribution between the as-grown and phosphorus diffused sisters at the bottom of the two ingots is small despite the lower total iron concentration. It was noted by Coletti *et al.* [21] that a higher crystal defect concentration existed at the bottom and top of their ingots relative to the middle portion of the ingot, possibly limiting the extraction of precipitated iron. Dislocations and their effects on gettering are not incorporated currently into the simulation, and thus the model poorly predicts the gettering response there.

In general, high concentrations of precipitated metals are seen here to persist throughout the phosphorus-diffusion process for a large range of iron concentrations. This result reflects the minimal adaption of phosphorus diffusion profiles in use today for multicrystalline silicon from those developed originally for higher-quality monocrystalline wafers, even though large concentrations of precipitated metals are known to exist in multicrystalline silicon wafers. Furthermore, anti-reflection coating deposition and the firing step, while critical to the functioning of a high-quality silicon solar cell, produce minimal changes in the precipitated iron distribution. Sim-

ilar results indicating little change in iron-silicide precipitates during extended low-temperature annealing have been found previously [42, 163].

4.5.4 Using As-Grown Iron-Silicide Precipitate Distribution Data to Predict Interstitial Iron Concentrations after Gettering

The permanence of iron precipitates during standard phosphorus diffusion has direct implications on final cell efficiency because the precipitated iron distribution determines the interstitial iron concentration after processing. Having measured the iron-silicide precipitate density and size distribution experimentally, we can collect the remaining inputs required for an I2E simulation of the gettering process on these ingots such that there are now no free metal-dependent parameters and compare such simulations with the experimental interstitial iron data. Accordingly, the as-grown total and Fe_i concentrations along the two ingots were taken or interpolated from published data [21, 26, 160].

Using the same gettering time-temperature profile as in Section 4.2, the simulated post-processed interstitial iron concentration is shown for the 53 ppmw and 200 ppmw ingots in Fig. 4-8(a) and 4-8(b) respectively. As-grown total iron concentrations are shown as full squares, as-grown Fe_i concentrations are shown as full triangles, and experimental post-processed Fe_i values [21, 164] are shown as open diamonds.

The post-processed Fe_i concentrations were simulated first at the three ingot heights investigated (full circles), using as an input the median as-grown precipitate radius measured by μ-XRF at that ingot height. When using the measured median radius at each height in addition to the interpolated as-grown total and interstitial iron concentration values, most of the simulated values are of the same order of magnitude as the experimental values, and some of them match almost exactly.

In particular, excellent agreement is seen between the simulated and the experimental values at the top of both ingots. Due to the distribution of precipitated iron that places them on border between Region 2 and Region 3 (from Fig. 4-1(a)),

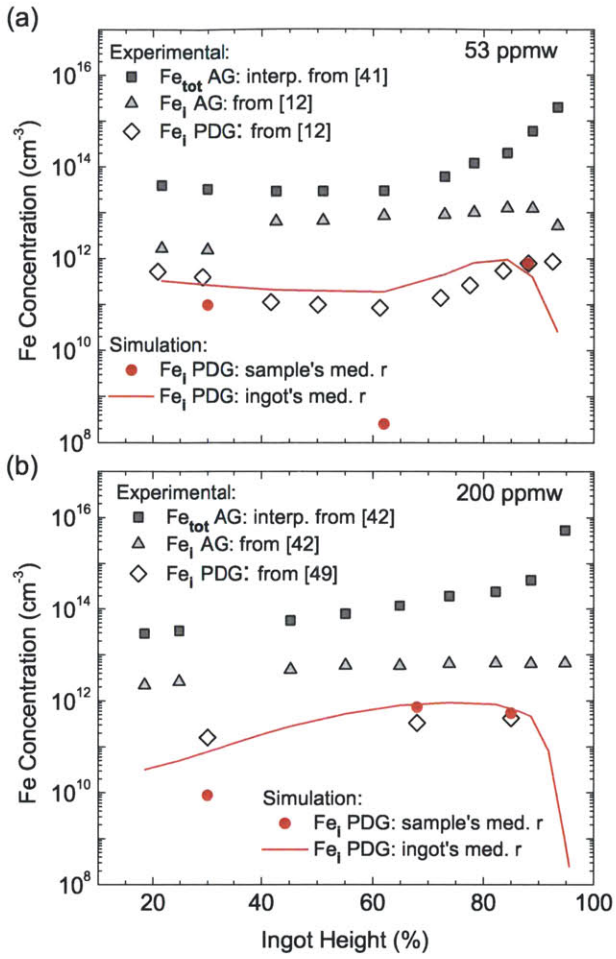

Figure 4-8: Shown here for (a) the 53 ppmw Fe ingot and (b) the 200 ppmw Fe ingot are the interpolated experimental values of the as-grown total Fe concentration (full squares) [26, 160], the as-grown Fe_i concentration (full triangles), and the post-gettering Fe_i concentration (open diamonds) [21, 160, 164]. The post-gettering Fe_i concentrations are also simulated: (full circles) at the three different ingot heights measured by µ-XRF using the median as-grown precipitate radius measured in the as-grown sample at that ingot height, and (solid line) assuming the as-grown radius everywhere in the ingot to be the median radius of all precipitates measured by µ-XRF in that ingot. Reproduced from [159].

the samples from the top of both ingots are predicted to have high interstitial iron concentrations after gettering. In fact, Coletti *et al.* measured high interstitial iron concentrations after gettering at the tops of these ingots (open diamonds in Fig. 4-8) and concluded that the high interstitial iron concentrations resulted from poor gettering in these top regions, ultimately leading to the degraded solar cell perfor-

mance seen there [21, 160]. The kinetics model predicts this poor gettering response directly from the as-grown distribution of precipitated iron measured at the top of these ingots.

Similarly, for the sample from the middle of the 200 ppmw ingot, a high interstitial iron concentration after gettering is predicted, in good agreement with the experimental value, due to the significant dissolution of precipitates (the iron distribution parameters place it toward the bottom of Region 2 in Fig. 4-1). In the end, solar cells produced from the middle of the 200 ppmw ingot were also iron-limited [160].

On the other hand, the interstitial iron concentration after gettering in the middle of the 53 ppmw ingot is predicted to be very low as seen in Fig 4-8. The experimental interstitial iron concentrations after gettering in the middle of the ingot from Ref. [12] approach 10^{11} atoms/cm^3, low enough to result in good solar cell efficiencies that were comparable to efficiencies on uncontaminated references. The simulated Fe_i concentration at the middle of the 53 ppmw ingot is far lower, however, than the experimental values. One reason for this discrepancy at in the middle of the 53 ppmw ingot is certainly that only one large $FeSi_2$ precipitate was found in the as-grown sample (see Fig. 4-3), resulting in poor statistics when calculating the median precipitate radius.

For a broader comparison to the experimental data, we simulated the post-gettering Fe_i concentration at all ingot heights (solid line), using the median value of all as-grown precipitate radii measured by μ-XRF in that ingot (bottom, middle, and top). Taking the median value for precipitates from all heights reduces the dependence on precipitate size along the ingot height in the simulation. A comparison between the simulation at all ingot heights and experiment shows that the simulated Fe_i concentrations reflect the trend of the experimental concentrations well. When using the median as-grown precipitate radius of the entire ingot, however, the simulated Fe_i concentration is often slightly above the experimental values, though well within the same order of magnitude.

At the very top of the ingot, our simulation suggests that the Fe_i concentration decreases sharply due to internal gettering of Fe_i to a high density of precipitates

during cooldown. Although we likely overestimate this internal gettering effect, as noted above in Section 4.2, the Fe_i data of Coletti *et al.* [21] shown here plateaus at high ingot heights, rather than continuing to increase if it were merely following the total iron concentration. In the Fe_i data that Kvande *et al.* [26] measured on samples from the same ingots used here, a true decrease in Fe_i at the top of the ingot is seen. Thus, the trend toward increasing internal gettering at the top of the ingot predicted by our model is also evidenced in the experimental data, though to a lesser extent.

In general, the good agreement between the experimental and simulated trends and values supports our hypothesis that gettering kinetics is governed by the as-grown distribution of iron *in addition to* the total iron concentration, as suggested by simulations in Section 4.2. For high total iron concentrations, the dissolution of $FeSi_2$ precipitates during standard phosphorus diffusion gettering strongly influences final interstitial iron concentration.

4.5.5 Overcoming High Iron Concentrations

We have demonstrated that a large quantity of iron is present as iron-silicide precipitates along grain boundaries in as-grown mc-Si wafers and that most of these precipitates remain after standard P diffusion gettering, especially in regions of high as-grown iron concentrations greater than 1×10^{14} atoms/cm^3. Carrier recombination at the precipitates themselves might limit minority carrier lifetime after PDG as indicated by simulations [11, 49] and recent experimental results [49]. Furthermore, because precipitated iron remains during and after gettering, the distribution of precipitated iron remains critical in determining the final interstitial iron concentration. The continuous dissolution of precipitates during high-temperature gettering limits the efficacy of iron interstitial reduction during phosphorus diffusion, leaving these lifetime-limiting defects in high concentration.

Additionally, with much of the total iron concentration remaining after gettering in precipitate form, any fractional dissolution of the precipitates iron during subsequent high-temperature processing, *e.g.*, metal contact firing, can have a strong effect on the interstitial iron concentration. Lifetime degradation after firing has been seen by

a number of authors [110, 158, 161]. With a significant concentration of precipitated iron remaining in the bulk after phosphorus diffusion, the injection of interstitial Fe from dissolving precipitates must be carefully avoided. Lelièvre *et al.*, for example, add a "low-temperature anneal" to the traditional firing profile to reduce the final interstitial iron concentration [158].

Advanced PDG with time-temperature profiles tailored to as-grown material properties may help improve material and, ultimately, device performance on wafers with high as-grown iron concentrations, *e.g.*, from the top and borders of the ingot. As first suggested by Plekhanov *et al.* [39] for aluminum gettering and recently confirmed experimentally for PDG [49, 165, 166], the reduction of interstitial *and* precipitated iron can be improved using a variable-temperature processing step: a short annealing step at high temperature allows for the enhanced dissolution of precipitates while a subsequent annealing at lower temperature drives dissolved iron toward the P-diffused gettering layer, leading to improved minority carrier lifetime.

The response to PDG has also been demonstrated to vary for different precipitate size distributions [167]: for Si ribbon materials, containing a higher density of smaller precipitates due to fast cooling after crystallization [19], the dissolution of precipitates is enhanced during high-temperature processing due to a larger interface area between precipitates and the Si matrix. This enhanced dissolution may lead to an enhanced poisoning of the silicon bulk during relatively short high-temperature gettering processes if the cooldown is not appropriately slow. But, the high-temperature also offers the opportunity for a faster reduction of the total iron concentration when applied along with a careful cooldown to room temperature to control the final interstitial iron concentration.

Predictive simulation tools offer an opportunity for an easy and fast optimization of time-temperature profiles for different starting materials. Ultimately, the combination of inline characterization and pre-sorting methods [32, 168] with customized solar cell processing may allow manufacturers to narrow the standard deviation of solar cell efficiencies while increasing average values.

4.6 Conclusion

Using synchrotron-based μ-XRF, we directly confirm the presence of large iron-silicide precipitates remaining after a standard phosphorus diffusion in samples containing as-grown total iron concentrations of $3 \times 10^{13} - 6 \times 10^{14}$ atoms/cm^3. While the ingots in this study were intentionally contaminated, levels of contamination within the range of this study (or higher) are often seen in regions affected by the crucible walls or toward the top of ingots, even when grown with electronic-grade Si. The survey of as-grown and gettered iron distributions from three positions within the ingot provides insight into the changes in iron distribution as a function of increasing iron concentration. Both precipitate radius and density increase with increasing iron concentration, with much of the increase in total iron concentration accommodated by increasing precipitate size.

Process modeling focused on the kinetics of precipitated iron in silicon readily predicts the limited efficacy of the standard gettering process, and agreement between experiment and simulation of iron gettering supports the argument that gettering kinetics are mostly governed by precipitated iron. With the gettering model supported by the experimental data, such models should have strong support for use in exploring pathways to improved material performance. While processing after phosphorus diffusion may strongly affect interstitial iron, it appears that the precipitated iron distribution remains largely unaffected, emphasizing the importance of a high-temperature, extended profile of phosphorus diffusion for control of precipitated impurities. In addition, with firm evidence that precipitated iron remains considerably unaffected by standard gettering, one must consider the consequences of downstream process steps such as metallization firing on precipitate dissolution.

While developing knowledge of iron's as-grown distribution for the defect engineering community, these experimental efforts also directly complement the I2E model, serving as initial conditions. These measurements of iron distribution in representative materials provide a platform of knowledge of iron distribution in wafers that can be used as inputs to our model by other researchers and process engineers.

Chapter 5

Investigating The Effects of a Low-Temperature Anneal in Border Material from an Industrial Ingot

In this chapter the I2E kinetics simulation in combination with μ-XRF are used to elucidate how in-line phosphorus diffusion and post-diffusion low-temperature annealing (LTA) determine iron distribution and device performance in standard industrial material. Focusing on samples from the degraded ingot border region, we find that during an industrial in-line phosphorus diffusion, bulk precipitate dissolution is incomplete. Post-diffusion low-temperature annealing is not found to alter appreciably the size or spatial distribution of $FeSi_2$ precipitates, although cell efficiency improves due to a decrease in iron interstitial concentration. Gettering simulations successfully model experiment results and suggest the efficacy of high- and low-temperature processing to reduce both precipitated and interstitial iron concentrations, respectively.

5.1 Materials and Methods

Three adjacent wafers, of thickness 250 μm and boron-doped resistivity 1 Ω-cm, were selected from the upper part of a corner brick of an industrial, directionally-solidified multicrystalline silicon (mc-Si) ingot used in previous LTA experiments [37]. After saw damage etching, one sample was set aside as-grown (AG). The other two wafers underwent a belt phosphorus diffusion gettering (PDG) process at 900°C for 6 min followed by a 10 min cooldown, producing a 45-50 Ω/□ emitter. After PDG, one sample underwent a low-temperature anneal of 60 min at 500°C with the phosphorus glass still present (P+LTA) . Samples of 20×20 mm^2 were cut from the AG, PDG, and P+LTA sister wafers from the deteriorated region influenced by the crucible wall, as evidenced by low local internal quantum efficiency [37]. The same grain boundary in all wafers was selected for μ-XRF analysis. Electron back-scatter diffraction measurements (not shown) indicate that the grain boundary selected is random-angle with a misorientation of 44.5°. Such random-angle grain boundaries exhibit high levels of impurity decoration [152].

Prior to μ-XRF, all samples were cleaned in a sequence of acetone, ethanol, and isopropanol, followed by cleaning with HCl and HF. μ-XRF measurements were obtained at the Advanced Photon Source beamline 2-ID-D at Argonne National Laboratory

using a 10 keV incident X-ray beam with 200 nm full-width half-maximum spot size, taking 325 nm steps. National Institute of Standards and Technology standard reference materials 1832 and 1833 were measured in the same detector geometry to convert fluorescence counts into iron concentrations. After measuring the Fe distribution in the AG sample with μ-XRF, a 5×5 mm piece containing the grain boundary was removed and annealed *in situ* in a high-temperature sample stage [169] to compare the effects of an LTA with no gettering layer present (AG+LTA).

5.2 μ-XRF results on As-Grown, Gettered, and Low-Temperature Annealed Samples

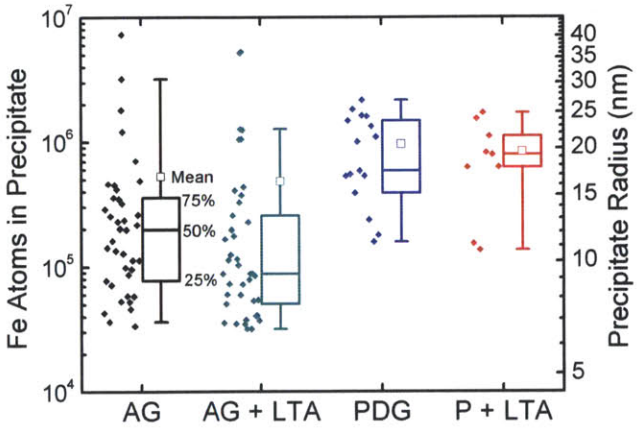

Figure 5-1: (Color online) Data points designate individual iron precipitates measured by μ-XRF. The number of atoms per precipitate is shown on the left and converted into an effective precipitate size on the right. The box plots to the right of the data for each sample show the 25th, 50th, and 75th percentiles of precipitate size in that sample. The open box shows the mean precipitate size. Low-temperature annealing is seen to have a minimal effect on the precipitated iron distribution. Reproduced from [42].

Iron precipitates were detected in all samples by μ-XRF, demonstrating that the thermal budget of the phosphorus diffusion was insufficient to dissolve and getter all metals. Of the other $3d$ transition metals, only copper was detected in the gettered samples by μ-XRF (not shown), but at concentrations roughly an order of magnitude lower than iron. Furthermore, copper is known to have a significantly lower impact on

lifetime than iron at the same concentrations [8, 24]. It is assumed in the simulations below that iron is the lifetime-limiting defect. Over 100 iron precipitates were analyzed in total. The number of atoms in each iron precipitate, shown in Fig. 5-1, was calculated from the iron loadings measured by μ-XRF assuming that all precipitates were at the surface. The effective precipitate radius, shown along the right axis of Fig. 5-1, was then calculated assuming a spherical precipitate of room-temperature $FeSi_2$.

Comparing the samples with LTA in Fig. 5-1 to the corresponding samples without LTA, we conclude that the 500°C LTA does not change the iron precipitate size distribution. Furthermore, comparing the AG and AG+LTA data, which show the μ-XRF analysis of the same precipitates before and after an *in situ* LTA, it is also clear that the LTA affects the density of precipitates minimally. However, there appears to be a decrease in precipitate density and increase in median precipitate size for the samples that underwent the 900°C P-diffusion process. Along the 35 μm length of the grain boundary observed in each sample, the AG sample contained 1.2 precipitates/μm, and the AG+LTA sample 1.3 precipitates/μm, while the PDG and P+LTA samples showed a decrease in precipitate density to 0.4 and 0.2 precipitates/μm respectively, with an effective 1/e attenuation length for the fluorescent Fe signal in our μ-XRF setup of 8 μm. If this decrease in density were due to coarsening, a broadening of the precipitate size distribution would be expected [170] but is not seen here. Further, recent results have shown internal gettering has much less of an effect than external gettering during annealing when surface sink layers are present [171, 172]. Therefore, the distribution changes after PDG seen in Fig. 5-1 suggest that dissolution due predominantly to external gettering eliminates smaller precipitates near the emitter. Although the two gettered samples both show the same trend with respect to the ungettered samples, we cannot eliminate the possibility that the variation may be due to wafer-to-wafer variations in local metal impurity concentrations [157].

5.3 Simulations Show a Low-Temperature Anneal Strongly Decreases Only the Interstitial Iron Concentration

To corroborate our experimental observations, we used our I2E diffusion-gettering simulator [103] to solve for the distribution of phosphorus and iron after processing and calculate an effective bulk lifetime. The lifetime and simulated phosphorus profile are then used as inputs for a PC1D device simulation [112]. We simulated the 6 min P-diffusion at 900°C followed by (a) a 10 min cool down to room temperature and (b) LTA at 500°C for 60 min. We chose an average as-grown precipitate radius of 17.5 nm, as measured by µ-XRF, and assumed an initial iron content for the corner wafer of 10^{15} cm^{-3}, typical of border region material [27, 123]. We assumed symmetrical gettering to both sides of the wafer. Fe-B pairs were taken to be completely dissociated, and iron precipitates were assumed to be uncharged, causing recombination of carriers at their surfaces. The simulation results for the processed Fe and effective lifetime distributions are shown in Fig. 5-2.

Due to iron segregation into the 0.3 µm emitter, the simulated bulk Fe content decreased by 35% with respect to the AG concentration for both processes. Beyond the emitter, both [Fe_{tot}] and [Fe_i] are depleted far into the bulk, as seen in Fig. 5-2(a), to depths much greater than the 8 µm information depth of the µ-XRF setup. Thus, returning to Fig. 5-1, the larger precipitate size seen in the PDG-processed samples could be a result of our high sensitivity to large precipitates in or near the heavily-segregated emitter, while the low measured density of precipitates in the PDG samples may reflect the shrinking of many precipitates below detection limits in the iron-poor region just beyond the emitter.

After the standard P-diffusion, the simulated [Fe_i], when averaged over the entire bulk, increased slightly with respect to the as-grown [Fe_i] due to bulk precipitate dissolution (increase seen at depths greater than 60 µm in Fig. 5-2(a)), which may explain the degradation of certain material types during high temperature process-

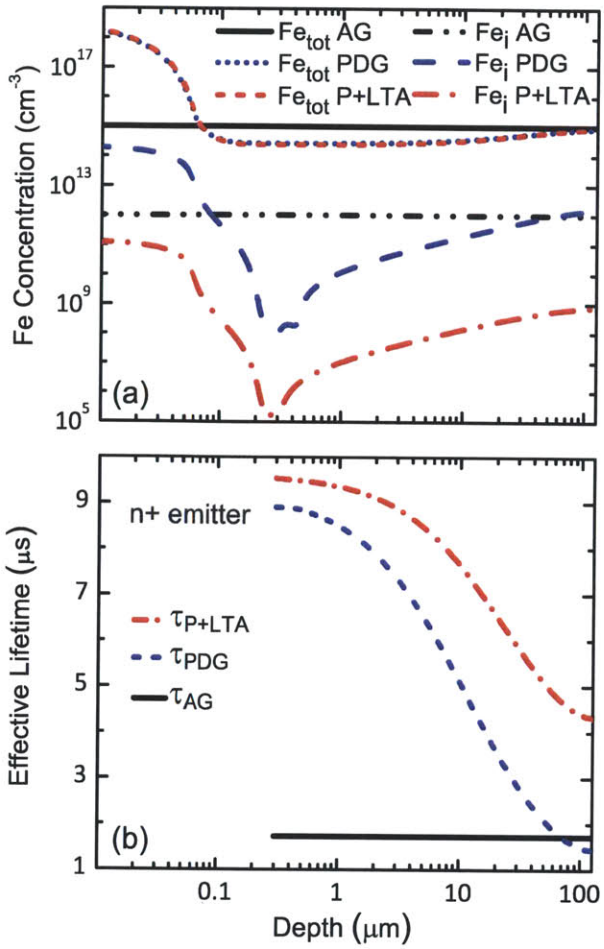

Figure 5-2: (Color online) (a) Simulated interstitial, Fe_i, and total, Fe_{tot}, iron profiles after standard PDG (blue) and after PDG followed by a 60 min LTA at 500°C (red). (b) The resulting effective lifetime in the p-type base for the two different processes is shown. After LTA, the bulk lifetime increases due to the large reduction of iron interstitial population.

ing [119]. When P-diffusion is followed by LTA, the $[Fe_i]$ decreased over three orders of magnitude. The reduction of $[Fe_i]$ during LTA led to an increase in bulk lifetime, as seen in Fig. 5-2(b). After LTA, the effective lifetime is dominated by the large amount of iron that remains as recombination-active precipitates.

5.4 Experimental Efficiency Improvements after Low-Temperature Annealing are Fully Attributable to the Decrease in Interstitial Iron Concentration

Using the simulated bulk lifetime and phosphorus profiles, we simulated a device in PC1D using typical parameters of a base contact of 0.8 Ω/cm^2, an emitter contact of 10^{-4} Ω/cm^2, a shunt resistance of 10^3 Ω/cm^2, a front surface recombination velocity of 10^5 cm/s, and a back surface recombination velocity of 10^3 cm/s. Simulation results for the two processes and the corresponding experimental efficiencies [37] are shown in Table 5.1. Excellent agreement is observed with the experimental data: $\geq 0.7\%$

Table 5.1: Experimental and simulated solar cell characteristics for the PDG and P + LTA processes

Treatment	PDG	P + LTA
Experimental efficiency(%)	**14.5**[a]	**15.2**[a]
Simulated efficiency (%)	**14.2**	**15.1**
Short circuit current density (mA/cm^2)	31.2	32.6
Open circuit voltage (mV)	594	604

[a] Ref. [37]

absolute higher solar cell efficiency was both simulated and measured for cells from border regions processed with LTA. The predicted change in the lifetime after LTA due to lower [Fe_i] fully explains the efficiency improvement that is observed experimentally. The combined modeling of lifetime improvement due to [Fe_i] reduction and experimental data on iron precipitates indicate that LTA acts on Fe_i with little effect on precipitated iron.

The removal of precipitated iron in such highly contaminated materials is severely kinetically-limited, since the bulk iron concentration is above the limit of solid solubility. Alternative processing will be required to further improve cell quality from ingot edge regions. To enhance precipitate dissolution, short plateaus at higher tem-

peratures must be considered, followed by an extended, variable-temperature cool down to drive segregation to the emitter in a process similar to that considered by Plekhanov *et al.* [39] and Seibt *et al.* [97].

5.5 Conclusion

To summarize, we observe that a standard phosphorus diffusion gettering process incompletely dissolves precipitated metals in mc-Si that contains high as-grown metal concentrations. The addition of a low-temperature anneal does not change the precipitated iron distribution significantly, as shown by μ-XRF analysis and simulations. Instead, simulations indicate that the improvement in device efficiency after a LTA can be fully explained by a sharp decrease in interstitial iron concentration. Subsequent modeling using I2E and experiments on low-temperature annealing have suggested that even a short extension of the traditional phosphorus diffusion profile can achieve much of the benefit in interstitial iron reduction [173]. An optimal annealing temperature for a given amount of extended annealing time can be chosen using the I2E model to minimize the final interstitial iron accounting for the tradeoff between lower diffusivity at lower temperatures but higher segregation driving force. With confirmation that the I2E model closely predicts the performance of experimental cells, the prospect of simulation-guided process optimization appears within reach for silicon solar cells.

Chapter 6

Improved Iron Gettering of Contaminated Multicrystalline Silicon by High-Temperature Phosphorus Diffusion

The efficacy of higher-temperature gettering processes in reducing precipitated iron concentrations is analyzed in this chapter by synchrotron-based micro-X-ray fluorescence. By measuring the same grain boundary before and after phosphorus diffusion in a set of wafers from adjacent ingot heights, the reduction in size of individual precipitates is measured as a function of gettering temperature in samples from the top of an ingot intentionally contaminated with iron in the melt. Compared to a baseline 820°C phosphorus diffusion, 870°C and 920°C diffusions result in a larger reduction in iron-silicide precipitate size. Minority carrier lifetimes measured on wafers from the same ingot heights processed with the same treatments show that the greater reduction in precipitated metals is associated with a strong increase in lifetime. In a sample contaminated with both copper and iron in the melt, significant iron gettering and dissolution of detectable copper precipitates is observed despite the higher total metal concentration. Finally, a homogenization pre-anneal in N_2 at 920°C followed by an 820°C phosphorus diffusion produces precipitate size reductions and lifetimes similar to an 870°C phosphorus diffusion without lowering the emitter sheet resistance.

6.1 Introduction

Novel processing schemes are needed to achieve high efficiencies in wafers coming from the "red zone" regions of the ingot, including the sides and top of cast ingot. One recently-developed technique, low-temperature annealing (LTA), the topic of Chapter 5, has been shown by a number of authors to reduce metal point defect concentrations, particularly iron, and improve lifetime in border region materials [33, 35–37, 60]. However, lifetimes after LTA often remain too low to meet todays efficiency requirements. LTA has limited effectiveness in part because precipitated metals remain inside the bulk [42]. In the bulk, precipitates can affect lifetime directly due to recombination at their surfaces [49, 171] or can re-dissolve during metallization firing, poisoning the bulk [110, 158, 161].

An alternative path to reduce the impact of metal impurities is to use higher-

temperature processes [39]. By increasing bulk solubility and speeding diffusion to near-surface gettering layers, higher-temperature diffusions are expected to improve the extraction of metals from the bulk, including deleterious iron. To date, experiments seeking lifetime improvements via high-temperature diffusions have seen mixed results. For example, Peters et al. [73], Ballif et al. [71], and Macdonald et al. [72] observed lifetime degradation after higher-temperature diffusions in a variety of silicon materials. In contrast, recent investigations by Osinniy et al. [165] showed lifetime improvement in solar-grade silicon after variable-temperature diffusion like the temperature profile recommended by Plekhanov et al. [39], supporting the hypothesis of improved impurity gettering indirectly through lifetime measurements. Similarly, Schön et al. [49] used a high-temperature pre-annealing in N_2 before phosphorus diffusion to improve lifetime. On the other hand, without showing lifetime improvement, Macdonald et al. [174] recently used micro-X-ray fluorescence measurements to determine that the percentage of surface area covered by iron-rich precipitates was lower after higher-temperature processing in wafers originating from the bottom of a cast ingot. While millimeter-sized areas were sampled, the collection of grain boundaries investigated in each sample differed as samples from adjacent ingot heights, *i.e.*, with the same grain structure, were not used.

In this contribution, we show direct evidence from synchrotron-based micro-X-ray fluorescence (μ-XRF) measurements that the precipitated iron concentration decreases significantly after higher-temperature phosphorus diffusion, using measurements of the same precipitates before and after phosphorus diffusion to draw conclusions. We then directly relate this decrease in precipitated metals to a marked improvement in bulk carrier lifetime in multicrystalline silicon contaminated with iron.

6.2 Materials and Methods

To determine the impact of gettering process conditions on metal precipitate distribution, we use μ-XRF measurements to characterize the metal distribution along the

same 20 μm length of a selected grain boundary before and after phosphorus diffusion gettering in a set of sister samples (samples from adjacent ingot heights). The phosphorus diffusion temperature profile was varied for each sister to investigate the effect of the gettering temperature.

6.2.1 Wafer and Region of Interest Selection

Full 15.6×15.6 cm^2 wafers were selected from approximately 92% (177 mm) ingot height from two ingots grown as part of the German SolarFocus project: one intentionally contaminated in the melt with 20 ppma Fe, and the other with 20 ppma Fe and 20 ppma Cu. Neutron activation analysis of the total impurity concentrations within these ingots as a function of ingot height can be found in Riepe et al. [22]. From this data, the Fe concentration in these wafers is estimated to be on the order of 10^{15} atoms/cm^3. In the Fe + Cu-contaminated samples, the Cu concentration should be on the order of 5×10^{16} atoms/cm^3, while the unintentional Cu concentration in the intentionally only Fe-contaminated samples should be 3×10^{14} atoms/cm^3. To focus on the impact of the intentional metal contamination and avoid the influence from impurity in-diffusion from the crucible wall, a stack of 5×5 cm^2 sister samples were cut from the full wafers from the region close to the center of the laboratory-scale ingot.

Electron-beam induced current and electron backscatter diffraction (not shown) were measured on one sister sample to select a recombination-active, random-angle grain boundary (GB) for μ-XRF investigation. This same GB was investigated in all samples except the sample that underwent a 920°C $POCl_3$ diffusion. In that sample, we investigated a second comparable, recombination-active, random-angle GB that was 2.5 cm from the GB of the other samples because no additional sister samples were available. Before as-grown μ-XRF measurement, the samples were etched with an acidic saw-damage etch followed by a cleaning sequence of hydrochloric acid (HCl) and hydrofluoric acid (HF).

6.2.2 Synchrotron-based μ-XRF Measurement Details

The μ-XRF measurements were conducted at Beamline 2-ID-D at the Advanced Photon Source at Argonne National Laboratory [147]. A 200 nm full-width half-maximum, 10 keV beam was used to map the sample using 220 nm steps, with a 1 second dwell time per pixel. NIST standard reference materials 1832 and 1833 were used to convert fluorescence counts into area concentrations.

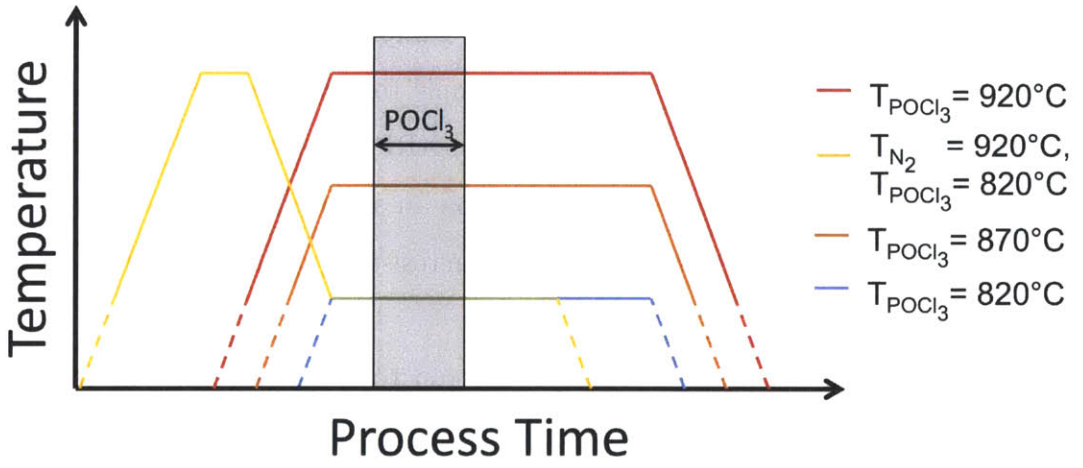

Figure 6-1: The time-temperature profiles of the phosphorus diffusion processing used in this study are shown. Three different peak temperatures were used: 820, 870, and 920°C. In addition, one sample was included that had a pre-anneal at 920°C, but the phosphorus diffusion was carried out at 820°C. Reproduced from [166].

To reveal the impact of the gettering step on the metal precipitate distribution, different phosphorus diffusions were conducted on the sister samples following synchrotron measurement of the as-grown metal distribution.

Three diffusions on the Fe-contaminated synchrotron samples used an extended drive-in step and varied in peak temperature only: at 820°C, 870°C, and 920°C. Total plateau time at process temperature was 76 min. The fourth Fe-contaminated synchrotron sample was first annealed in N_2 at 920°C before the temperature was lowered to 820°C for phosphorus diffusion and a standard drive-in (54 min total plateau time), inspired by lifetime improvements shown previously using pre-annealing [49, 161]. Ad-

ditionally, one Fe + Cu-contaminated sample was investigated at the synchrotron. It was processed using the 870°C phosphorus diffusion with extended drive-in, which gave the highest lifetimes after the extended gettering. All samples were pulled from the diffusion furnace at 820°C, and the cooling rate to 820°Cfor the higher-temperature processes was 6°C/min. The four different phosphorus diffusion time-temperature profiles used are shown schematically in Fig. 6-1.

Before high-temperature treatment, the synchrotron samples were cleaned again by a sequence of HCl and HF. Following the phosphorus diffusion, the phosphorus glass was removed from the surface of the sample by dipping in HF, but the phosphorus-diffused emitter was left on the sample to compare precipitates near to the surface that were observed in the as-grown μ-XRF measurements. The resulting sheet resistance of each process was measured on an additional boron-doped *p*-type Czochralski (Cz) monocrystalline wafer (1-3 Ω-cm resistivity) included in the diffusion for that purpose.

The samples were returned to the synchrotron to re-measure the metal distributions after phosphorus diffusion using identical experimental conditions to the as-grown measurements. Test samples produced by e-beam lithography containing metallic copper dots down to 8 nm in radius [150] were measured at the start of both synchrotron runs to verify that XRF sensitivity was indeed similar between runs.

6.2.3 Carrier Lifetime Measurements

Microwave photoconductance decay (μ-PCD) lifetime measurements were performed using a Semilab WT-2000 on a neighboring stack of 5×5 cm^2 sister samples (also unaffected by the crucible wall) to evaluate the impact of processing on high metal contamination levels.

Two sister wafers in the stack were measured as-grown. These as-grown lifetime samples were prepared using the same saw-damage etch and HCl and HF sequence as above. The remaining sisters in the stack, corresponding to the same heights as the synchrotron samples, were phosphorus diffused after saw-damage etch and HCl and HF cleaning according to the same profiles as described above (see Fig. 6-

1). These sister wafers, from lower ingot height to higher (centered around 92% ingot height), were assigned to the 820°C, 870°C, 920°C, and 920°C/820°C processes, respectively. Post-diffusion lifetime measurements were performed after etching back approximately 5-10 μm per side to remove the emitter using a fresh damage etch solution, followed by the same HCl, HF sequence before lifetime measurement.

All lifetime measurements reported here used a solution of iodine ethanol (I-E) to passivate the sample surfaces. To ensure comparability, all samples were additionally cleaned by piranha etch followed by an HF dip, and measured immediately following exposure to the I-E [175].

6.3 Results

6.3.1 Fe-contaminated synchrotron samples

The resulting iron concentration distribution maps from μ-XRF measurement of the Fe contaminated samples are shown in Fig. 6-2 according to the phosphorus-diffusion treatment that was applied to the sample. All maps in Fig. 6-2 are shown on the same logarithmic iron concentration scale to allow for direct comparison. We did not observe a significant amount of other metals by μ-XRF in the Fe-contaminated samples.

Precipitates the highly localized, high-concentration agglomerations of iron in the μ-XRF maps are seen in high density both before and after phosphorus diffusion. Many of the larger precipitates (e.g., those numbered for the 920°C PDG sample) are readily identifiable as the same precipitate before and after phosphorus diffusion, indicating that even the highest temperature phosphorus diffusion here (920°C) does not fully dissolve the largest precipitates. On the other hand, close inspection of the maps reveals that in each sample, a number of precipitates observed in the as-grown sample are not observed after phosphorus diffusion (e.g., those in bottom left and top left of 920°C PDG sample). The disappearance of some precipitates signifies their dissolution at least to levels below the detection limit of the measurement (estimated

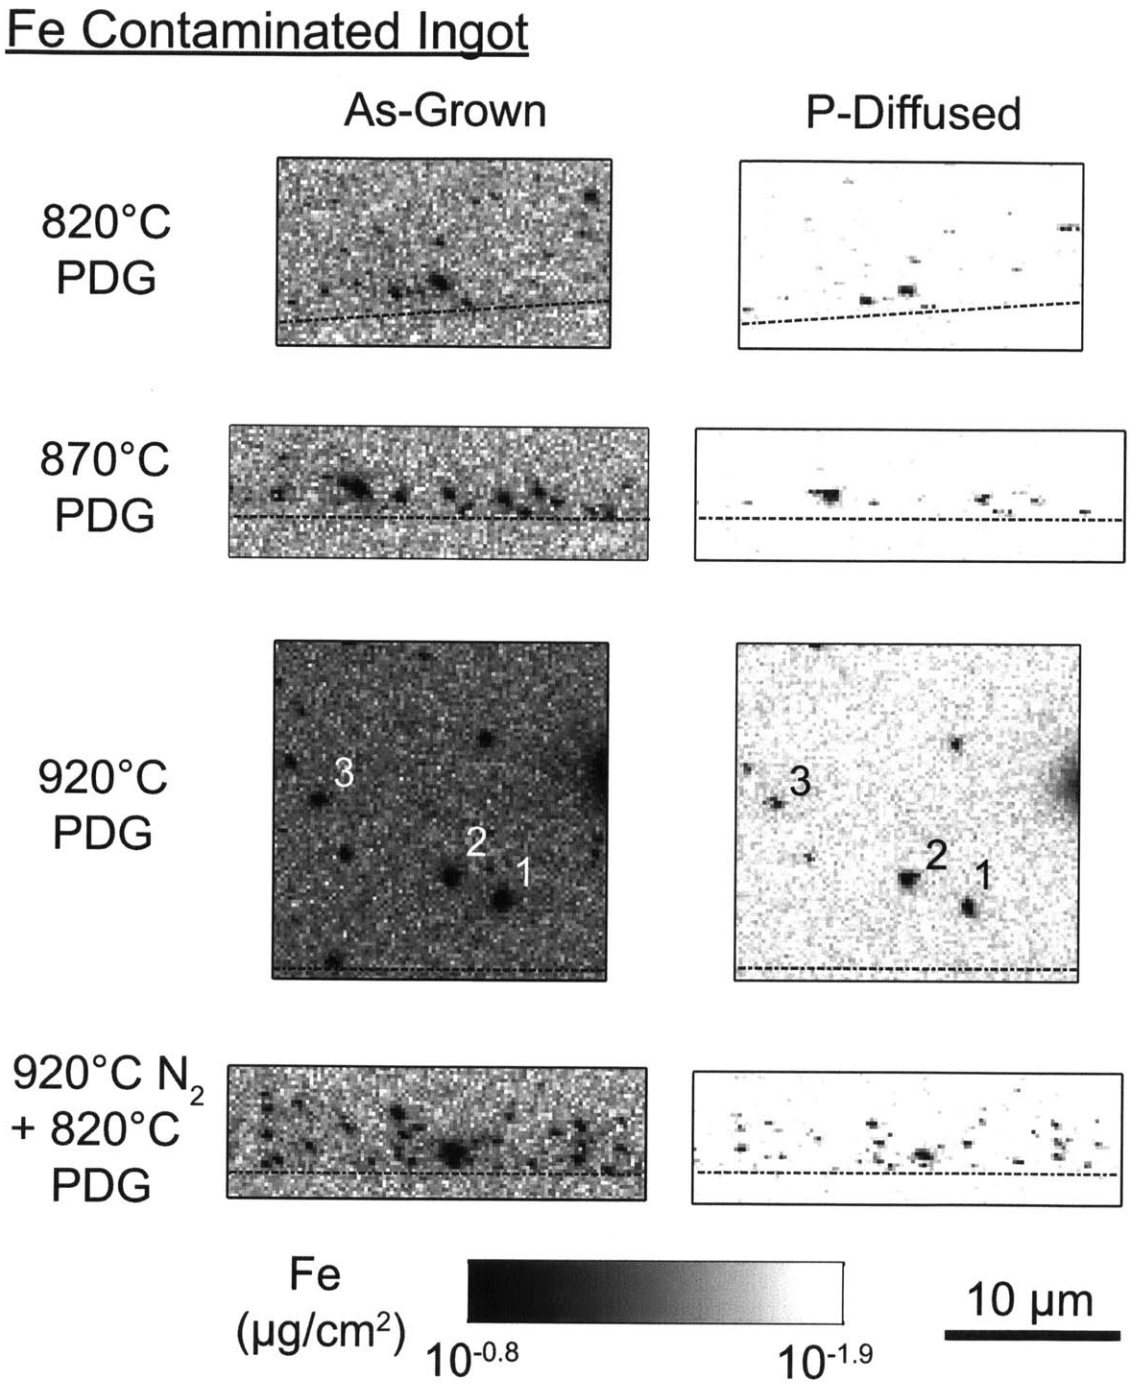

Figure 6-2: Maps of iron concentrations ($\mu g/cm^2$) in each sample as-grown (AG) and after phosphorus diffusion gettering (PDG) as measured by μ-XRF. The same grain boundary was measured for the 820°C, 870°C, and 920°C N_2 + 820°C PDG samples. A dashed line indicates the position of the grain boundary in each image. Reproduced from [166].

to be 5×10^4 atoms/precipitate). Thus, the extended, high-temperature phosphorus diffusion gettering steps appear to fully dissolve some precipitates while the largest remain.

Inspection of linescans of the as-grown maps (not shown) reveals that the background level increases in the regions immediately surrounding precipitates. In addition, the 920°C PDG sample has the largest as-grown precipitates and also the highest background level. While some contribution to the background is attributable to impurities in the detector and scattering off instruments in the beamline hutch, the relationship between the presence of large precipitates and a higher background level suggests the background is also indicative of the presence of high iron concentration in the sample. Thus, a decrease in background level after phosphorus diffusion indicates a decrease in iron concentration in the region measured. From the difference in the coloring of the background level between the as-grown maps and the phosphorus-diffused maps in Fig. 6-2, it appears that in all cases the phosphorus diffusion step reduces the level of iron in these samples.

To further quantify the changes in distribution of iron-rich precipitates in each sample, the size of each precipitate in Fig. 6-2 was estimated using an automated routine to identify iron precipitates from pixels containing iron concentrations more than three standard deviations above the average background level. An effective precipitate radius was calculated from the concentration maps assuming all iron atoms within high-Fe pixels could be attributed to a single spherical β-$FeSi_2$ precipitate located at the surface of the sample. Similarly, for the Fe + Cu-contaminated sample, high-Cu pixels were attributed to Cu_3Si precipitates. Further details on the μ-XRF data analysis and the extraction of precipitate distributions can be found in Chapter 3.

The resulting precipitate size distributions are shown for each sample in the as-grown state (AG) and after phosphorus diffusion gettering (PDG) in Fig. 6-3. In some cases, the PDG μ-XRF map was larger with respect to the AG map in order to be sure to capture the same region of interest for the direct comparison above. In an effort to include all available data, in Fig. 6-3 all measured precipitates are shown, including those from the PDG maps outside the region shown in Fig. 6-2.

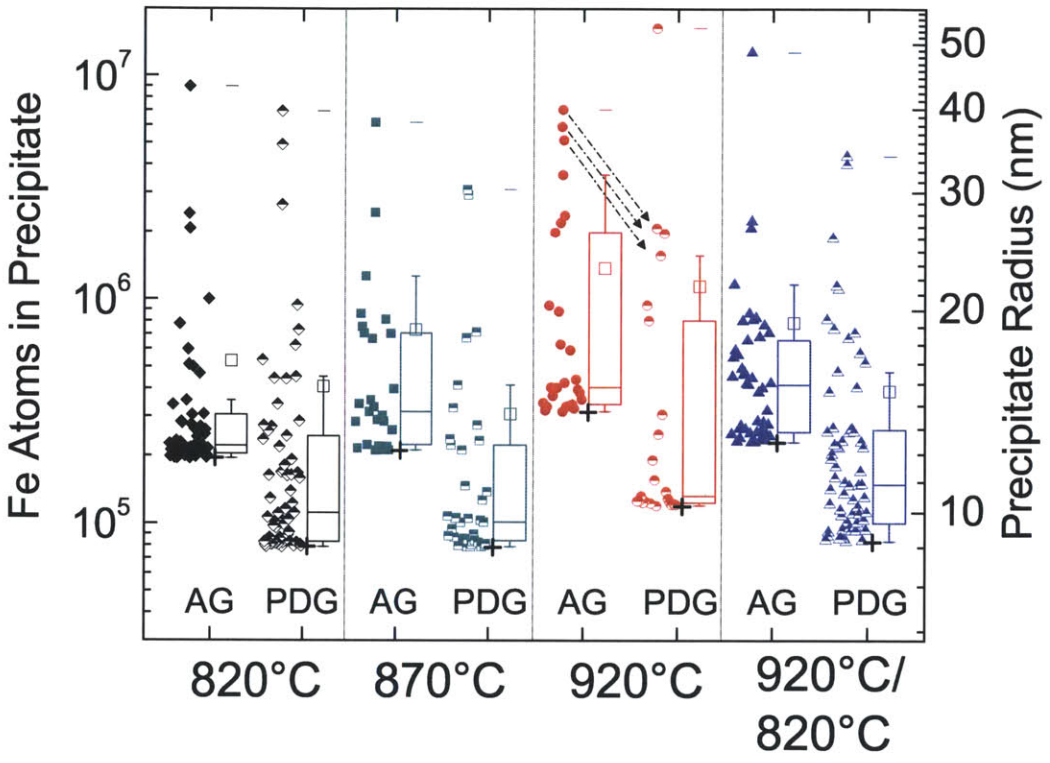

Figure 6-3: Quantified precipitate size distributions for the Fe contaminated samples in the as-grown state and after gettering according to the diffusion process used for each sample. The boxplots to the right of each data set show the quartiles of the distribution, while the open squares show the mean precipitate size. The cross at the bottom of each distribution indicates the effective detection limit in each sample below which precipitates could not be separated from the background. Reproduced from [166].

The shift toward smaller precipitate sizes after phosphorus diffusion is clear in Fig. 6-3, where each point represents one precipitate. In each sample, all distribution parameters are smaller after phosphorus diffusion relative to the as-grown state, as revealed by comparison of the distribution boxplots to the right of the data.

It is important to note that the cross at the bottom of each distribution in Fig. 6-2 represents the limit beyond which precipitates could not be distinguished from the background level of the map. Due to higher levels of iron measured in the as-grown samples, the background level is higher, perhaps excluding existing smaller precipitates from inclusion in the distribution parameter estimates. However, as mentioned above, many of the precipitates can be clearly, uniquely identified before and after diffusion such that small precipitates omitted from the as-grown distributions should be similarly excluded from the phosphorus-diffused distribution analysis.

Table 6.1: Measured linear density of Fe precipitates along the GB of interest for each sample in the as-grown state and after phosphorus diffusion gettering.

	Fe 820°C $POCl_3$	Fe 870°C $POCl_3$	Fe 920°C $POCl_3$	Fe 920°C/ N_2 820°C$POCl_3$	Fe + Cu 870°C $POCl_3$
As-Grown Pcp. Density (μm^{-1})	2.3	0.99	1.05	1.81	1.66
PDG Pcp. Density (μm^{-1})	1.55	0.95	0.95	1.61	1.37

In addition to the size distribution, the linear density of precipitates at each GB can be extracted from the μ-XRF maps before and after gettering and is shown in Table 6.1 for all samples in this study. In addition to a reduction in precipitate sizes, the density of precipitates decreases in all cases after PDG.

Three particles in the 920°C sample, labeled 1, 2, and 3 in order from largest to smallest in Fig. 6-2, are among the many precipitates identified as the same particle before and after gettering. Their sizes before and after gettering are shown using arrows to guide the eye in Fig. 6-3. From the analysis of Fig. 6-3, it is seen that all three precipitates decrease in size dramatically, on a scale approaching an order of magnitude.

To support this automated analysis drawn from the slightly undersampled 220 nm pixel maps of Fig. 6-2 (full width at half maximum, FWHM, 200 nm), 100 nm pixel maps of these three particles were taken and are shown in Fig. 6-4. In the oversampled (100 nm step size) particle maps in Fig. 6-4, the peak concentrations measured at each particle (listed in Table 6.2) are generally higher than in the larger-area maps sampled at 220 nm, readily expected since the precipitates are uniformly smaller than the beam FWHM. The trend from these oversampled maps is consistent with the analysis of Fig. 6-3, which uses slightly undersampled data.

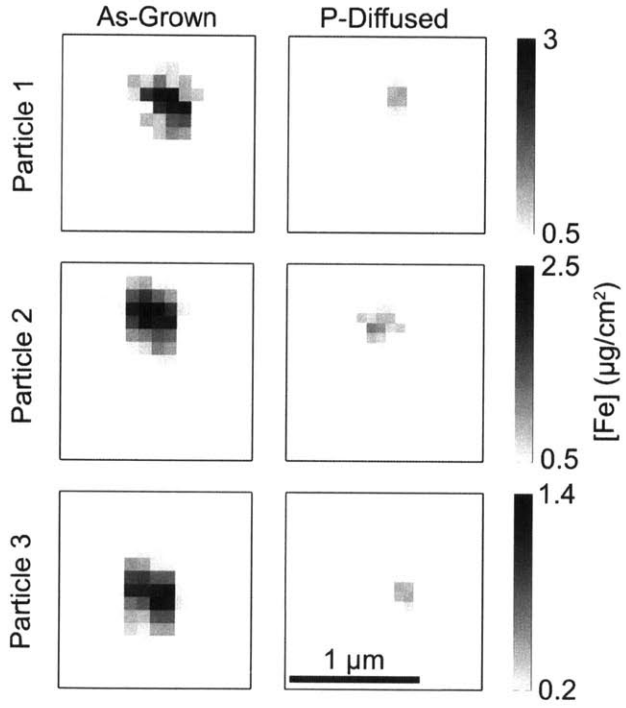

Figure 6-4: High-resolution maps (100 nm pixel size) of the three particles, highlighted in Fig. 6-2 and Fig. 6-3, from the 920°C sample before and after gettering. The number of iron atoms in each precipitate is reduced by at least one order of magnitude. Reproduced from [166].

Despite the undersampling for the data of Fig. 6-3, the trend of decreased precipitate size after phosphorus diffusion is uniformly consistent across all precipitates, even though the measurements were from two separate synchrotron runs. The random repositioning of the < 50 nm precipitate within the 200 nm FWHM beam spot during the post-gettering measurement should lead to a run-to-run intensity varia-

Table 6.2: The peak concentrations at the three particles in the 920°C PDG sample measured using 100 nm and 220 nm step sizes are shown before and after gettering.

	Sampling	Particle 1 ($\mu g/cm^2$)	Particle 2 ($\mu g/cm^2$)	Particle 3 ($\mu g/cm^2$)
As-Grown	220 nm	2.1	1.7	1.5
	100 nm	3.1	2.5	1.4
P-Diffused	220 nm	0.6	0.6	0.5
	100 nm	1.1	1.2	0.5

tion of a particular precipitate of approximately a factor of 2. As can be seen from Table 6.2, the concentrations do in fact agree within the expected factor of 2 when moving between 100 nm and 220 nm sampling. Yet, no case of precipitate size increase after diffusion is observed, indicating that we are most likely underestimating the overall precipitate size reduction, by up to a factor of 2. This underestimation should be uniform in all samples due to its random nature. In short, Fig. 6-3 should be viewed as a conservative estimate of iron reduction.

Because not all precipitates could be investigated at the highest resolution due to simple experimental time constraints, a comparison of precipitate size before and after gettering was performed using the data of Fig. 6-3 for all precipitates that were definitively identified from Fig. 6-2 before and after gettering. The results of this precipitate-by-precipitate size reduction analysis are shown in Fig. 6-5 in terms of the percent reduction observed for the three Fe-contaminated samples processed at different diffusion temperatures (820°C, 870°C, and 920°C). Each point in the figure represents the reduction in size of a single precipitate after gettering. Notably, the median precipitate size reduction increases for higher temperatures, as can be seen from the boxplots to the right of the data for each sample which again show the quartiles of the distribution. While the increase in median size reduction with increasing temperature is small, and a strong reduction is seen for each temperature, the 920°C sample also exhibits no precipitates with $\lesssim 60\%$ size reduction.

Almost all of the precipitates in the fourth Fe synchrotron sample, with the N_2 pre-anneal at 920°C, were impossible to identify definitively as the same precipitate

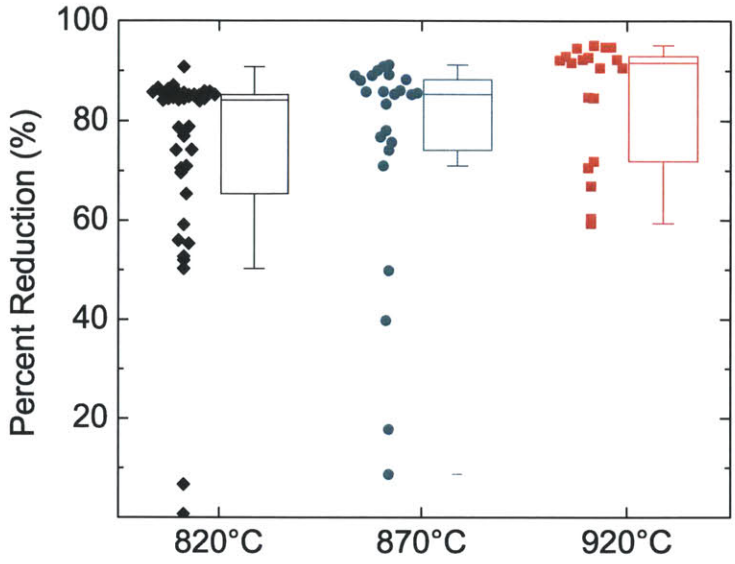

Figure 6-5: The percent reduction in number of precipitated atoms of all precipitates identifiable before and after phosphorus diffusion is shown for the three different temperature processes. A trend toward higher median reduction with higher gettering temperature is seen. Reproduced from [166].

before and after gettering, mostly due to small changes in their relative positions and sizes. While preventing the same precipitate reduction analysis for the sample as in Fig. 6-5, this difficulty in establishing congruency in the precipitates before and after high-temperature processing is in accordance with the hypothesis that a high-temperature pre-annealing step before gettering can be useful for homogenization.

To establish whether the differences in precipitate size reduction seen in Fig. 6-5 are meaningful, pair-wise non-parametric rank-sum tests were conducted [176]. The resulting p-values are reported in Table 6.3. As can be seen from the table, the difference between the 920°C and 820°C distributions is highly significant, while the comparison between the 920°C and 870°C samples is also significant at the 0.05 level. The difference in precipitate size reduction between the 870°C and 820°C samples is significant only at the 0.065 level. Note that the p-values reported have no multiple-comparison adjustments, though a typical Holm-Bonferroni correction

Table 6.3: The median reductions for the three different gettering temperatures for the Fe samples are listed. In addition, the p-values resulting from a rank sum test of different medians between the samples are listed. The reduction during a 920°Cdiffusion is significantly different than lower temperatures.

	Fe 820°C $POCl_3$	Fe 870°C $POCl_3$	Fe 920°C $POCl_3$
Median % Reduction	84.1	85.3	91.5
Rank-Sum Test with 820°Cprocess	–	$p = 0.064$	$p = 1.3\times10^{-3}$
Rank-Sum Test with 870°Cprocess	–	–	$p = 0.018$

does not change the conclusions of the tests [177]. Thus, we conclude that the higher-temperature gettering processes employed here reduce iron precipitate sizes more than the baseline 820°C process.

6.3.2 Fe + Cu contaminated synchrotron sample

In addition to the Fe-contaminated samples above, a single Fe + Cu contaminated sample, processed using an 870°C phosphorus diffusion, was measured as an initial investigation into co-contamination effects. The resulting μ-XRF maps of the Fe and Cu distributions are shown on the same concentration scale before and after gettering in Fig. 6-6(a).

In the as-grown maps, one can see that all Fe and Cu precipitates are co-located and appear generally smaller than those seen in the Fe-only samples. In addition, in the Cu as-grown map, the position of the grain boundary is indicated by the line running across the map leading to contrast between the upper 80% of the map and the bottom 20%. A small amount of the elastically scattered beam affects the Cu map because the 10 keV beam is close in energy to the Cu K_α radiation at 8047.78 eV. The contrast stems from a difference in elastic scattering between the two grains within the limited solid angle of the detector, appearing here due to a favorable positioning of the sample with respect to the detector.

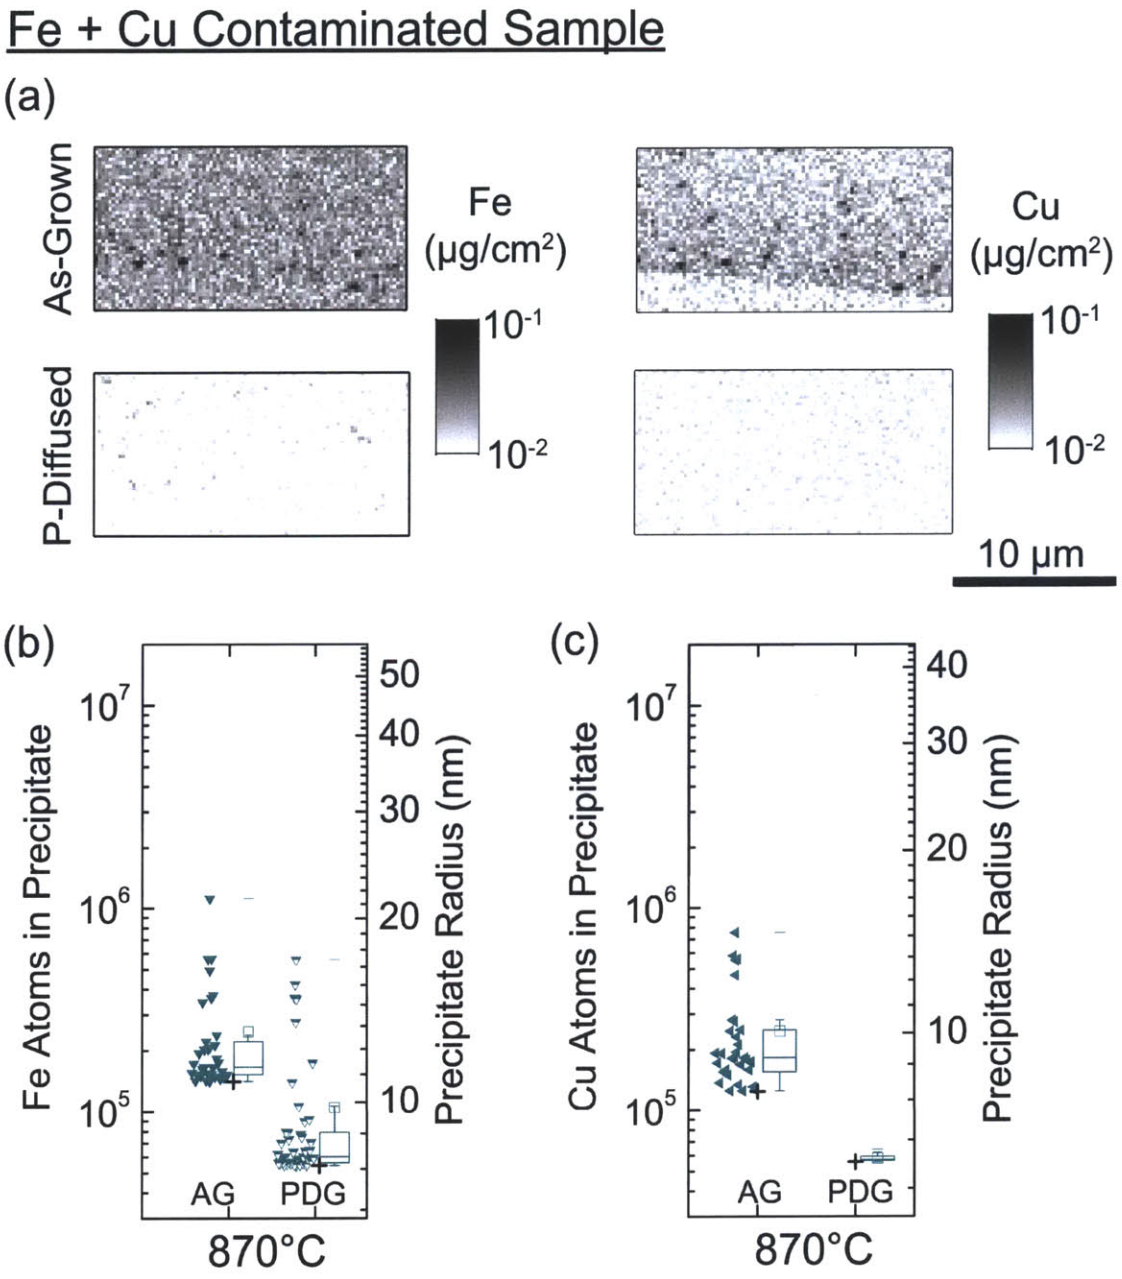

Figure 6-6: μ-XRF maps of Fe and Cu concentrations in the co-contaminated sample are shown in (a). The precipitate size distributions corresponding to these maps are shown for (b) Fe and (c) Cu. No copper-rich precipitates are observed above the background after gettering, while small iron-rich precipitates remain. The higher background level seen for Cu after gettering than for Fe is presumably due to the proximity of the Cu edge to the energy of the incident beam (10 keV) rather than high, uniform Cu concentration. Reproduced from [166].

After phosphorus diffusion, some faint Fe-rich precipitates remain. Note that the concentration scale is shifted toward slightly smaller values than in Fig. 6-2 in order to image these faint Fe precipitates. While some Fe remains, no Cu precipitates are immediately observable after gettering. Although the background level in the Cu XRF map after gettering is higher than for Fe, it is expected that this background level reflects some bleed-in from the 10 keV elastically-scattered beam rather than a uniform Cu concentration (which would be above 10^{16} atom/cm^3 and would limit the lifetime to < 1 μs [178]).

Fig. 6-6(b) contains the quantified Fe precipitate distribution extracted from the as-grown and gettered μ-XRF maps, while Fig. 6-6(c) shows the corresponding Cu precipitate distribution. The as-grown Fe distribution in the Fe + Cu sample is indeed smaller with respect to those in the Fe-only samples, ranging from about 10^5-10^6 atoms/precipitate in the Fe + Cu sample as compared with about $2\times10^5 - 10^7$ atoms/precipitate in the Fe-only samples. Interestingly, after gettering, the background level for Fe is also particularly low in the Fe + Cu sample, allowing for detection of smaller precipitates than in the Fe-only samples. Again, the lower background level may be a secondary indication of effective removal of total iron concentration. The automated analysis also reveals no copper-rich precipitates above the background level after gettering.

6.3.3 Lifetime Measurements

To test the hypothesis that precipitated metal reduction is associated with improved carrier lifetime, lifetime measurements were taken on a stack of sister wafers neighboring those studied by μ-XRF and processed with the same diffusion profiles.

The area-averaged lifetime results for the sister samples after phosphorus diffusion are summarized in Table 6.4. As-grown lifetimes in the iron-limited wafers are below 3 μs, which is also typical of material originating from the red zone of ingots. Both the arithmetic average and the harmonic average are included, as the harmonic average is a better assessment of future solar cell quality than the simple arithmetic average [179].

Table 6.4: Average μ-PCD [μs] lifetime results measured on sister samples from the three materials, processed by different phosphorus diffusions. The arithmetic and harmonic average lifetimes are calculated to weight the heterogeneity of the distribution. Also included are the sheet resistance measurements made on the Cz samples included in each of the diffusions.

		P-diffusion Process			
Ingot	Measurement	820°C	870°C	920°C	920/820°C
Fe	AG	0.68 ± 0.4			
	PDG_{arith}	10.3 ± 2.9	18.9 ± 7.2	17.9 ± 10.5	16.0 ± 5.3
	PDG_{harm}	9.4	16.3	13.1	14.2
Fe + Cu	AG	1.2 ± 0.5			
	PDG_{arith}	25.5 ± 15.8	32.7 ± 22.0	32.2 ± 25.3	33.1 ± 22.5
	PDG_{harm}	17.1	19.6	16.6	19.6
Cz	AG	38.7 ± 5.0			
	PDG_{arith}	392.5 ± 91.5	179.3 ± 26.4	131.3 ± 18.1	230 ± 56.8
	PDG_{harm}	374.4	175.5	129.2	219.3
Sheet Resistance (Ω/sq.)		62.6 ± 1.6	23.2 ± 0.6	10.6 ± 0.2	67.6 ± 1.7

Lifetimes improve dramatically moving from the the 820°C diffusion to either the 870°C and 920°C diffusion temperatures, but variability also increases when considering the 920°C diffusion with respect to the 870°C diffusion. The lifetime average after the 920°C N_2/820°C PDG gettering step shows less variation in the Fe-contaminated samples, supporting the idea of the pre-anneal as a homogenization step, though the average lifetimes of the higher-temperature phosphorus diffusions are not achieved. Further results and discussion of the spatially-resolved lifetime distribution can be found in [166].

6.4 Discussion

6.4.1 Iron-Contaminated Material

The principal hypothesis of this work is that higher processing temperatures offer a path to improve material quality in iron-limited materials through increased extraction of precipitated metals during the phosphorus diffusion gettering process. By accelerating the external gettering process, higher-temperature processing should improve bulk lifetime and allow for acceptable device efficiency even from degraded as-grown materials. Attacking total metal concentrations can also be vital to device performance outside of the impact on lifetime. Precipitates have shown to lead to early breakdown and shunting, which must be avoided for long module lifetimes [54, 55].

The improved gettering of precipitated metals with increasing processing temperature seen herein with synchrotron-based μ-XRF measurements supports the stated hypothesis. Using a set of sister samples subjected to increasing processing temperatures and measured before and after gettering, we confirm on the microscopic scale that increasing temperature results in enhanced iron precipitate size reduction (Fig. 6-2-Fig. 6-5). The largest median precipitate size reduction of over 90% occurs in the sample gettered at 920°C. In correspondence with this improved precipitate size reduction at higher temperatures, nearly 80% higher average arithmetic lifetimes are seen in both the 870°C and 920°C Fe-contaminated samples, though the final harmonic lifetime is highest after 870°C gettering.

The largest relative lifetime improvements by moving to higher temperatures are seen in Fe-contaminated samples. In fact, the absolute improvement in lifetime between the higher temperature gettering steps and the baseline 820°C gettering process (8.6 and 7.6 μs for the 920°C and 870°C P-diffusions, respectively) is almost as large as the difference between the 820°C gettered lifetime and the as-grown lifetime (9.6 μs).

In general, when cooling from higher-temperature anneals, care must be taken to avoid poisoning the wafer with iron point defects. Fast cools have been shown to lock

in high concentrations of iron point defects, sharply reducing lifetime. I2E simulations for higher gettering temperatues readily predict this effect. The interstitial iron concentration after gettering is shown in Fig. 6-7 for gettering temperatures ranging from 820 to 870°C considering 3 different cooldowns: (1) pulling immediately from the plateau temperature with free-cooling to room temperature (red squares), (2) cooling to 800°C within the furnace prior to pulling the boat out (blue triangles), and (3) cooling to 750°C prior to pulling (black circles). In-furnace cooling was taken to be a typical 3.33 °C/min. With no slow cooling, higher interstitial iron concentrations are

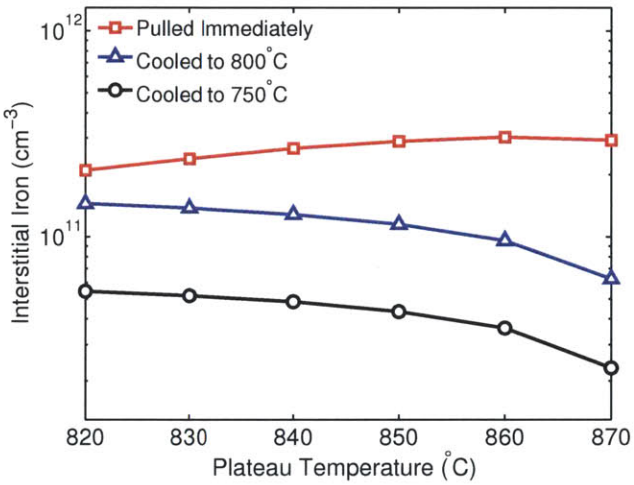

Figure 6-7: The effect of a high-temperature gettering temperature is coupled to the cooldown from high temperature. If wafers are immediately pulled, simulations predict high interstitial iron concentrations remaining after the process which can lead to low lifetimes. Even a limited amount of slow-cooling can produce large reductions in the post-process interstitial iron concentration.

seen after the application of higher gettering temperatures. In industrial production, it is common practice to pull wafers from the furnace at high temperatures, effectively quenching in high-temperature impurity distributions. A modified cooling treatment may thus be necessary to adapt higher-temperature anneals in practice, and to avoid low lifetimes reported in previous studies that investigated higher-temperature gettering [71, 73]. Rapid thermal processing [73] and its fast ramping rates in particular may be problematic. On the other hand, Fig. 6-7 shows that slow-cooling for just 15 min to 800°C has the potential starkly reduce final interstitial iron concentration, and the effect is greater for diffusions conducted at higher temperatures because more

of the precipitated iron has been extracted. Further cooling to 750°C for example continues to reduce interstitial iron, though other defects may begin to dominate lifetime once very low interstitial iron concentrations have been achieved.

In the above experiments, process time has been held constant to isolate the effect of temperature. Other possible constraints, such as the sheet resistance of the emitter, have been relaxed, although all diffusions here achieved a surface phosphorus concentration above the electrically-active limit of $\approx 3 \times 10^20$ atoms/cm^3 [78] which should mean that both segregation and relaxation gettering mechanisms are active in all samples [94].

In general, a lower sheet resistance does not necessarily imply better gettering. This is in part because the emitter sheet resistance is dominated by the depth of the high-concentration plateau in the kink-and-tail profile [180]. Iron profiles are generally seen to be confined to a region very near the surface that is much narrower than the high-concentration plateau even for extended gettering steps that result in lower sheet resistance (see e.g. [99]). Rather than strictly depending on the emitter sheet resistance, gettering efficacy depends strongly on the exact nature of the time-temperature profile.

Extending gettering time and increasing gettering temperature are two ways to achieve a longer diffusion length of iron during processing, but an increase in temperature rather than time is more powerful because of the Arrhenius relationship of iron diffusivity. Iron solubility also has an Arrhenius form, and so precipitate dissolution should be markedly higher with higher temperatures. Demonstrating the impact of the kinetics, we have shown that for approximately the same process throughput we can reduce the precipitated iron concentration after higher-temperature gettering, and we correlate this reduction with improved lifetime.

A heavier phosphorus diffusion is produced, but the etch back of a strongly-diffused emitter has already been achieved for selective-emitter designs [125–127]. Such processes could potentially be applied to produce uniformly better blue response for an industrial emitter after a heavy diffusion aimed at gettering precipitated iron.

6.4.2 Co-contamination effects (Fe + Cu)

In the Fe + Cu co-contaminated synchrotron sample, the as-grown distribution of precipitates is shifted toward smaller sizes. Previously, mixed-metal silicides have been observed [181, 182], but limited solubility of one element in the others silicide often results in two co-located precipitates [22]. Such precipitates are predominantly a single metal, justifying our approximation of independent spherical precipitates of each metal type for analyzing size distributions in the Fe + Cu μ-XRF map. Still, interactions between the precipitating species appear important in explaining the smaller measured as-grown distribution. Copper has a large energy barrier to nucleation in silicon [183], but with co-precipitation a reduction in this barrier may explain the more widespread precipitation seen here in the Fe + Cu sample relative to the Fe only samples.

The change in the iron distribution after 870°C PDG is also more significant in the Fe + Cu sample than in the Fe-only contaminated sample, as there is no overlap in the distribution parameters before and after gettering in the Fe + Cu sample (Fig. 6-6). Two effects may explain this more significant extraction of precipitated metals. First, for a given gettering temperature, smaller precipitate sizes should result in improved gettering due to faster kinetics [167]. Second, reversing the apparently enhanced co-precipitation process, during gettering as the solubility rises with temperature dissolving Cu_3Si precipitates may create a high local concentration of vacancies, speeding $FeSi_2$ precipitate dissolution.

Ultimately, higher post-gettered lifetimes are measured in the Fe + Cu samples compared to the Fe contaminated samples irrespective of the process, despite higher initial total metal concentrations in the Fe + Cu wafers. However, the relative lifetime increase resulting from high-temperature annealing appears less impressive in the Fe + Cu materials tested, which have smaller (or undetectable) iron-rich precipitates. In such co-contaminated material, higher-temperature annealing appears to have less of a beneficial effect on lifetime improvement, since lower-temperature $POCl_3$ appears to satisfactorily remediate precipitated iron.

6.4.3 Homogenization Pre-annealing

As a possible alternative to higher P-diffusion temperatures in metal-contaminated samples that maintains a sheet resistance appropriate for devices, a pre-anneal in nitrogen at temperatures higher than a subsequent P-diffusion is seen here to improve lifetime over a baseline process, confirming previous results [49, 161]. Little internal or external gettering is observed for high-iron wafers processed solely in an N_2 environment [11], as anneals in N_2 should have no strong driving force for a unidirectional flux to the surface. Instead, largely isotropic dissolution/diffusion should occur at process temperature, resulting potentially in homogenization. Along these lines, in the pre-annealed sample here, the position of the precipitates after the pre-annealing and diffusion is similar to but not exactly the same as the as-grown distribution. A similar pattern of dissolution and reprecipitation at almost identical, but slightly differing sites was seen by *in-situ* μ-XRF measurements during heating and cooling of an Ni-contaminated Cz sample with precipitation occurring at bulk microdefects [184]. In addition to the slightly differing precipitate positions in the pre-annealed sample, from our μ-XRF measurements we conclude that it appears there is a significant decrease in total metal concentration during the two-step 920°C N_2/820°C P-diffusion process. The changes in the precipitated iron distribution after two-step treatment appears closest to the corresponding results from the 870°C P-diffusion and improved over the baseline 820°C P-diffusion.

6.5 Conclusion

In this contribution, synchrotron-based μ-XRF provides a microscopic assessment of the effect of gettering temperature on precipitate size reduction. We analyze the precipitate size distribution in the 2D spatial maps to quantify process comparisons statistically. By rigorously measuring the same precipitates before and after high temperature processing, we confirm that higher-temperature diffusions successfully extract a greater fraction of the precipitated metal content. Lifetime mappings on a macroscopic scale of neighboring samples relate this reduction in precipitated metal

content to significant improved lifetime in iron-dominated samples.

When considering the design of a gettering process it is critical to assess what defect limits lifetime in that material. If the limiting defect is iron, precipitate dissolution and bulk iron removal increases markedly moving to higher temperatures, demonstrated directly using μ-XRF measurements herein. Notably, when moving to higher processing temperatures, the interstitial iron concentration may be relatively higher at the start of the cooldown from diffusion temperature due to the higher solubility. This high interstitial concentration requires further defect engineering during cooling and subsequent metallization, such as slow-cooling or low-temperature annealing. Nonetheless, higher processing temperatures may be one of the few ways to circumvent higher contamination levels using process control to achieve dramatically improved lifetimes that enable quality cell efficiencies.

Homogenization pre-annealing may also prove a useful industrial process near-term since little is required for implementation, though benefits may be limited in comparison with high-temperature P-diffusion. A combination of pre-annealing and moderately elevated phosphorus diffusion may prove most beneficial. A large parameter space remains open here for exploration of such combined process optimization.

Chapter 7

Extracting Design Guidelines for Phosphorus Diffusion Gettering

7.1 Introduction

In silicon photovoltaics, we are inspired by the remarkable scaling laws of the semiconductor industry. This scaling was partly enabled by the ability to predict defect evolution during device processing, derived from deep knowledge of device-limiting defects and their kinetics.

In this chapter, I present a structured approach to accelerated, simulation-guided defect engineering that we use to optimize the critical phosphorus diffusion step with an emphasis on achieving successful gettering of metal impurities. First, using the I2E simulation tool, I perform a sensitivity study detailing the response of a typical solar silicon wafer to changes in the time-temperature profile of the phosphorus diffusion gettering process. Having evaluated the sensitivity to process variables *in silico*, I present detailed experiments examining the dependence of the device-limiting interstitial iron concentration on time-temperature process variables predicted to be most and least sensitive, demonstrating the key physical underpinnings of material and process performance. Then, I extend the capability of our simulation to solving the inverse design problem, using a genetic algorithm to rapidly explore the process design space in an effort to move toward complete material-specific process optimization. Finally, I extract design principles for phosphorus diffusion gettering based on the learnings from simulation and experiment.

7.2 Analyzing the Sensitivity of Interstitial Iron Concentration to the Time-Temperature Profile of Phosphorus Diffusion

The I2E simulation tool is used to identify the key time and temperature process variables that impact phosphorus diffusion gettering of interstitial iron, focusing on interstitial iron because of its large lifetime impact. Six variables that define the time-temperature parameter space are chosen for investigation, two each from the

loading, the hold, and the unloading steps of the annealing process. They are: the temperature at which the samples are loaded into the furnace (Load Temp.) and the rate at which they are heated to process temperature (Heat Rate), the temperature of the diffusion step (Plateau Temp.) and its duration (Plateau Time), the cooling rate down from the process temperature (Cool Rate) and the temperature to which the samples are cooled within the furnace before being unloaded (Unload Temp).

To understand the relative impacts of each of these variables, a one-factor-at-a-time approach is taken where the variable values are varied around a base case to high and low values. The base case and the high and low values of each variable tested are shown in Table 7.1. The range of each variable is chosen so that any change in a variable is equivalent to a 15 min deviation in total process time with respect to the base case. For example, changing the Cool Temp from the base case value of 750°C to 700°C takes 15 min more of cooling at the base case cooling rate of 3.3°C/min. Similarly, increasing the plateau temperature from the base case of 845°C to 870°C without changing any other variable takes 15 extra min of heating and cooling time at the base case rates of 3.33°C/min. In this way, the various time-temperature profile variables can be compared on an equal basis with respect to their impact on manufacturing throughput.

Variable	Low Value	Base Case	High Value
Load Temp. (°C)	700	750	800
Heat Rate (°C/min)	2.2	3.3	7
Plateau Temp. (°C)	820	845	870
Plateau Time (min)	15	30	45
Cooling Rate (°C/min)	2.2	3.33	7
Unload Temp. (°C)	700	750	800

Table 7.1: The values used in the sensitivity analysis of the time-temperature profile variables are shown. Each parameter was varied, one at a time, to a high or low value, while the other variables were kept at the base case. Each change in a variable corresponds to a 15 min change in total process time.

The interstitial iron concentration after gettering is plotted as an indicator of process sensitivity to the variable for each case in Fig. 7-1. For an assumed initial total iron concentration of 10^{15} atoms/cm^3 distributed at 25 nm precipitates and

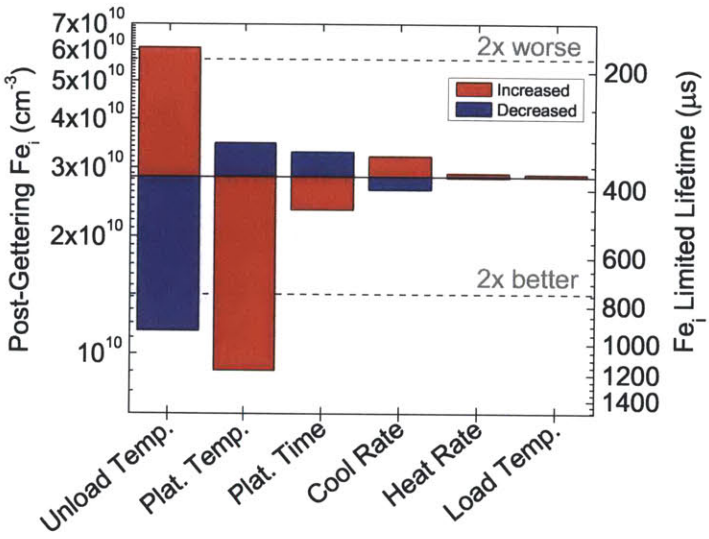

Figure 7-1: The sensitivity of the post-gettering interstitial iron concentration is plotted for a 15 min deviation from a base case for 6 time-temperature variables. If the variable value increases (decreases), it is designated as high (low) in the figure. Cool Temp, the temperature at which the samples are cooled to before unloading, and the Plateau Temperature are seen to have the largest impact on final interstitial iron concentration.

initial interstitial iron concentration of 10^{11} atoms/cm^3, the base case profile results in a final [Fe$_i$] just less than 3×10^{10} atoms/cm^3. The process variables are arranged from left to right in order of decreasing impact on outcome. Reducing the Unload Temp and increasing the Plateau Temp lead to significant reductions in interstitial iron. Conversely, increasing the Unload Temp significantly increases interstitial iron concentration at the end of the process, more than doubling it. Increasing the plateau time leads to a small reduction in interstitial iron, while the cool and load rates and the load temperature have very little impact.

7.3 One Factor at A Time: Load Temperature, Unload Temperature, and Plateau Time

To verify these processing predictions, the effect of the load temperature, the unload temperature, and the duration of the high-temperature plateau are explored further

in simulation and experimentally. From the sensitivity results in Fig. 7-1, the unload temperature is expected to have a large impact of final interstitial iron concentration, while extending the high-temperature plateau time should lead to some marginal decrease in the interstitial iron concentration. The effect of the load temperature is also explored here as an experimental control, since no change in final interstitial iron concentration is expected. Of the variables analyzed in Fig. 7-1 that are not investigated below, increasing the plateau temperature has been shown to lead to improved gettering results in Chapter 6. The cool rate and the heat rate are not explored here, though their impact is expected to be small, particularly the impact of a change in heat rate.

7.3.1 Simulated Effect of Load Temperature for Different Plateau Times

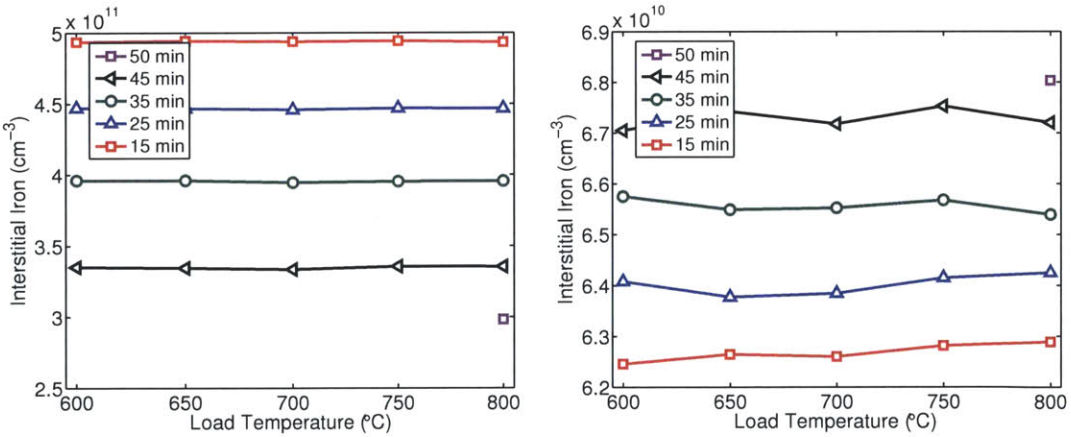

(a) Initial Total Iron: 5×10^{13} atoms/cm^3 (b) Initial Total Iron: 5×10^{15} atoms/cm^3

Figure 7-2: Simulated results of the effect of loading at low temperature and extending the high-temperature plateau time for high- and low-iron cases. No significant change in interstitial iron is noted for changes in loading temperature; a slight change is observed for variations in plateau time, consistent with Fig. 7-1.

First, a more detailed simulation of the dependence of final interstitial iron concentration on the load temperature, unload temperature, and plateau time is completed to provide finer-grained analysis than the base, low, and high case analysis in

Fig. 7-1. Accordingly, the post-gettered interstitial iron concentration is plotted as the load temperature is varied from 600 to 800°C and the plateau time is varied from 15 to 45 min in Fig. 7-2. A low total iron case (5×10^{13} atoms/cm^3) and a high total iron case (5×10^{15} atoms/cm^3) are shown in Fig. 7-2(a) and 7-2(b), corresponding to approximate total iron concentrations for monocrystalline silicon materials and heavily-contaminated multicrystalline materials, respectively. As can be seen quite clearly, the load temperature has no significant impact on the gettering performance. The plateau time, on the other hand, does show some effect, especially with lower [Fe_i]. The change in interstitial iron after a longer plateau time is predicted to be larger for the lower total iron case, a readily understandable result since the precipitated iron concentration that determines the kinetics of the gettering process is much lower than for the high iron case. For the low iron case, a decrease of about 30% in the interstitial iron concentration is seen when doubling from the standard 25 min diffusion to a 50 min diffusion. The high iron case actually shows a small increase in interstitial iron concentration of $< 10\%$ when plateau time is doubled. As precipitates shrink during gettering, their dissolution accelerates, which contributes to the higher interstitial iron concentration after gettering in the high total iron case. In the low total iron case, precipitates near the gettering layer dissolve completely, and so no longer inject interstitials into the bulk, allowing for reduced interstitial iron concentration after gettering. The simulated ramp rate up to the process temperature was 10°C/min (rather than the base case of 3.33°C/min in Section 7.2) to match the capabilities of our laboratory furnace for experimental verification below – the change in load rate had no effect on simulations. Also included in Fig. 7-2 is a 50 min gettering process with a load at 800°C, corresponding to the extended process tested experimentally in the next section.

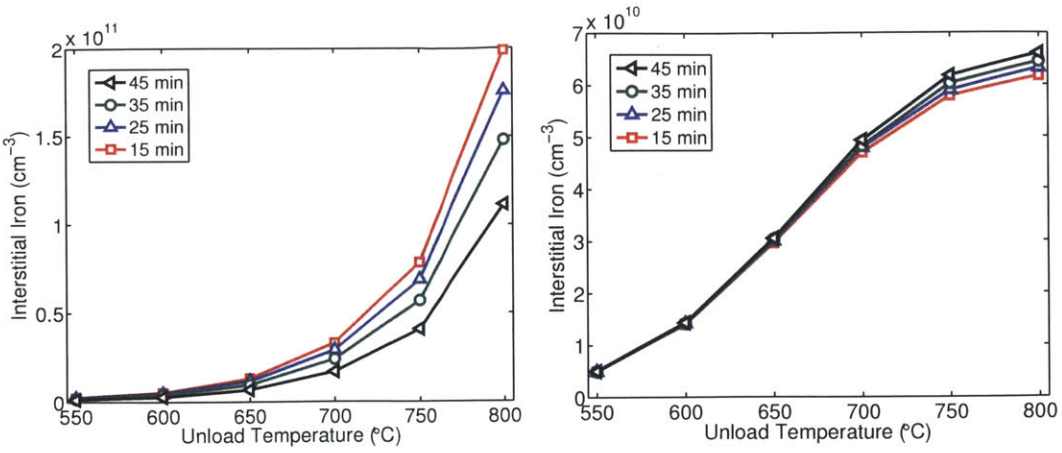

(a) Initial Total Iron: 5×10^{13} atoms/cm^3 (b) Initial Total Iron: 5×10^{15} atoms/cm^3

Figure 7-3: Simulated results of the effect of unloading the wafers from the furnace at low temperature. Also shown for comparison is the impact of extending the high-temperature plateau time.

7.3.2 Simulated Effect of Unload Temperature for Different Plateau Times

The dependence of the post-gettering [Fe_i] on the temperature that the wafers are cooled to is explored for the same low and high total iron concentrations in Fig. 7-3. For a relative comparison, the effect of extending the plateau time is again included. One can see that the unload temperature is the strongest predictor of interstitial iron concentration after gettering for both cases, as seen in the sensitivity analysis above. The interstitial iron concentrations fall sharply when cooling to lower temperatures before unloading. The simulated ramp rate down to the unload temperature was 3-4°C/min, again matching the experimental capabilities of a typical diffusion furnace.

7.3.3 Experimental Parameter Variation: Load Temperature, Unload Temperature, and Plateau Time

Two typical materials, one multicrystalline (mc-Si) and one monocrystalline (sc-Si), were chosen for an experimental investigation of the impact of these time-temperature variables. A set of 175 μm thick p-type multicrystalline silicon wafers (0.7 Ω-cm

resistivity) from adjacent ingot heights were selected from 90% total ingot height from a cast industrial ingot grown from recycled croppings of previous ingots. The recycling of the metal-rich tops of the ingots results in higher contamination levels. Mass spectrometry measurements indicate these wafers have iron concentrations of $\gtrsim 10^{15}$ atoms/cm^3. For the monocrystalline wafers, a set of standard 290 μm thick p-type CZ wafers (4 Ω-cm resistivity) was used. Monocrystalline wafers typically have iron contamination levels at or below mass-spectroscopy limits of about 1 ppb, or 5×10^{13} atoms/cm^3.

Sample Preparation and Annealing Parameters

The mc-Si wafers were etched in a CP4 acid solution (15 HNO_3:5 CH_3COOH:2 HF) to remove saw damage, followed by a KOH dip to remove re-plated porous silicon left by the CP4 etching. All wafers were then RCA cleaned [4]. Following cleaning, a 20 nm Al_2O_3 surface passivation layer was deposited by ALD (Cambridge Nanotech), followed by a 400°C anneal for 10 min in N_2 to achieve good surface passivation [185, 186]. Lifetime measurements of float-zone silicon passivated at the same time as the experimental samples achieved lifetimes of over 1 ms, indicating excellent surface passivation quality ($\lesssim$ 10 cm/s). QSSPC was used to take interstitial iron measurements on these as-grown wafers. The first measurements were taken within 30 s after exposure to 10 flashes of a 10 sun lamp to break Fe-B pairs and at 30 s intervals thereafter for 3 min. Lifetime measurements were then collected after 2-4 hours to allow for repairing. Calculation of interstitial iron concentrations was conducted at an injection condition of 10^{16} cm^{-3}, far away from the crossover point for the Fe-B system [137].

After as-grown lifetime measurements, phosphorus diffusions were conducted in a $POCl_3$ furnace (Tystar Corporation). Four different phosphorus diffusion time-temperature profiles were used:

- a **standard** process where wafers were introduced to the furnace at 800°C, ramped to 845°Cfor a 25 min P-diffusion, followed by an immediate pull with no cooling within the furnace itself

- a **load** process where wafers were loaded into the furnace at 600°C, followed by a ramp to 845°Cfor a 25 min P-diffusion and immediately pulled, as in the standard process
- an **extended** process which matched the standard process except the plateau time was doubled to 50 min
- an **unload** process which matched the standard process except that after the 25 min hold at the plateau, the wafers were slowly-cooled at 3-4°C/min within the furnace until 550°C before unloading

These four time-temperature profiles are shown schematically in Fig. 7-4(a). To each diffusion profile, four mutli- and four monocrystalline samples were assigned to evaluate the gettering effect of the process. The ramp rate up to process temperature was 10°C/min in all cases. After phosphorus diffusion, the phosphosilicate glass was stripped off in an HF solution. A CP4 etch was used to remove the emitter, followed by a KOH dip, and an RCA clean. ALD was then used to re-apply a Al_2O_3 passivation layer, as above. Post-diffusion interstitial iron measurements followed the same protocol as the as-grown measurements.

Process Variation Results

The average interstitial iron concentration after gettering is shown for the mono samples in Fig. 7-4(b) and for the multi samples in Fig. 7-4(d). Focusing first on the mono samples, the standard process leads to a final [Fe_i] of roughly 10^{11} atoms/cm^3. Within the error of the experiment, indicated in the figure by the standard deviation of the four samples at each condition, the load process resulted in the same final interstitial iron concentration. The extended gettering process led to a small decrease in interstitial iron. A significant decrease in interstitial iron concentration is seen by treating the samples with the unload process, which cools down to 550°C before pulling.

The experimental results from the multi samples show many of the same trends, as seen in Fig. 7-4(d). The load process does not yield significantly different results than

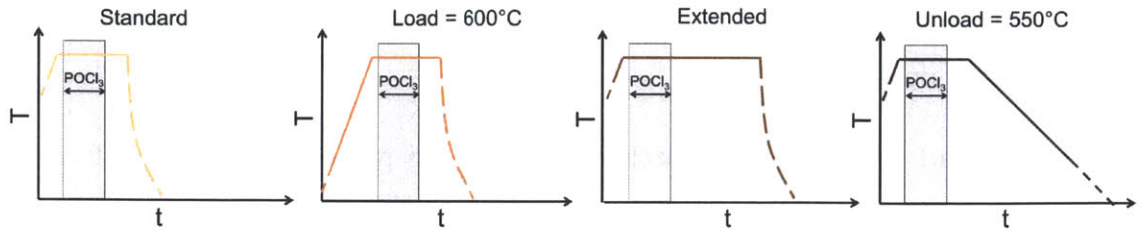

(a) Time-Temperature Profile Variations

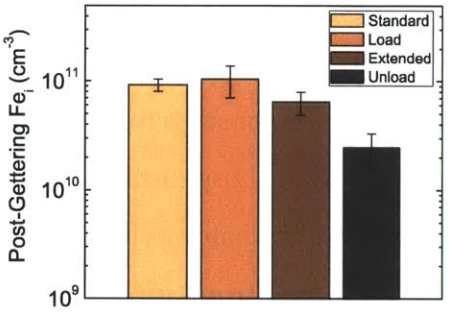

(b) sc-Si: measured Fe_i concentration

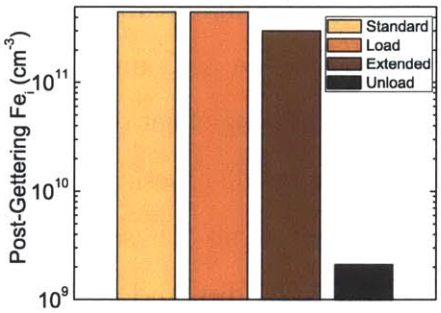

(c) sc-Si: simulated Fe_i concentration

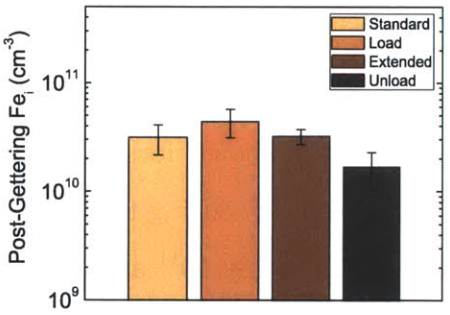

(d) mc-Si: measured Fe_i concentration

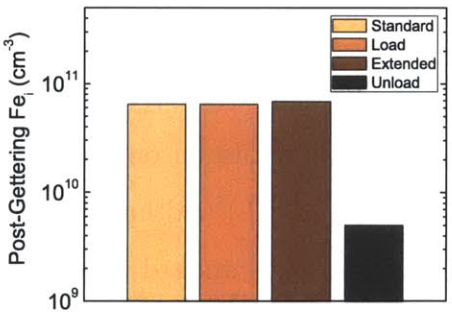

(e) mc-Si: simulated Fe_i concentration

Figure 7-4: Simulated and experimental gettering results compared for one-factor-at-a-time experiments, for monocrystalline silicon (b,c) and multicrystalline silicon (d,e). Time-temperature profiles corresponding to each condition are shown in (a).

the standard process. But, nor does the extended gettering yield lower interstitial iron, in contrast to the mono results. The unload process still leads to a reduction in interstitial iron concentration after gettering, approaching the detection limit of the measurement at 10^{10} atoms/cm^3. Of note, for the multi samples the average interstitial iron concentration after gettering was uniformly lower than that for the mono samples for all processes (*e.g.* the standard process was just over 3×10^{10} atoms/cm^3 for the multi as compared to nearly 10^{11} atoms/cm^3 for the mono), most likely as a result of internal gettering to structural defects.

Simulations using the same temperature profiles as were used in the experiment can help us understand these results and the differences between the response of the low total iron mono wafers and the high total iron multi wafers. A total concentration of iron for the mono wafers of 5×10^{13} atoms/cm^3 and 5×10^{15} atoms/cm^3 for the multi was assumed for the simulations. Fig. 7-4(c) shows the simulations for the mono wafers that correspond to the experimental measurements of Fig. 7-4(b). Likewise, Fig. 7-4(e) shows the simulated results for the multi samples of Fig. 7-4(d).

The general trends of the experiment agree very well with the simulations. First, the load process does not differ from the standard process significantly in either material in both the experiments or simulations. For the extended process, in the mono samples we see an improvement from an extended gettering experimentally that is predicted by simulation. In contrast, for the multi samples, no improvement is seen with an extended gettering step experimentally, but this lack of improvement is also predicted by simulation. With higher total iron concentration, the gettering kinetics are severely limited by precipitated iron in the multi samples (see Section 4.2.1), and doubling the gettering time has little effect. Instead, higher temperatures might be used to good effect as seen in Chapter 6, relying on the Arrhenius relationship of solubility and diffusivity with temperature.

There are several possible reasons why the simulation overpredicts the reduction in interstitial iron during the unload process in both the multi and mono. First, the solubility of iron below 800°C has been shown by Murphy and Falster [108] to be dependent upon thermal history, perhaps introducing unexpected effects outside of a

simple Arrhenius relationship. Second, the rate constant used as a fitting parameter to develop the effective segregation coefficient that drives iron gettering has been largely developed for higher temperatures [98], though using updated values including low-temperature data [99] does not affect the simulation results significantly. The general trend though, that in both high and low iron cases slow-cooling with a low unload temperature reduces the interstitial iron concentration dramatically is well confirmed between the simulation and experiment.

The lower interstitial iron concentrations for the standard, load, and extended processes in the multi samples as compared with the mono samples is also reproduced by the simulation. The density of precipitates is much higher in the simulation of the multi sample due to the higher total iron concentration, and internal gettering during the cooldown from process temperature leads to a lower [Fe_i] for multi samples with respect to the mono samples. While absolute values are not exactly reproduced by the simulation in all cases, that all of the above trends in the data are in good agreement between the experimental and simulated results provides a high degree of confidence in generalizing these result across silicon solar cell materials and further validates the I2E model.

7.4 A Genetic Algorithm to Tailor and Optimize Phosphorus Diffusion

The ultimate question of interest here is what time-temperature profile is optimal for processing a wafer containing a given starting concentration and distribution of iron. With the fast simulation run time of the I2E simulation (1-2 min typical), the flexibility exists to use *in silico* optimization methods rather than relying on extensive designed experiments to answer this question. To move away from the forward problem, where the time-temperature profile is used as an input to the I2E simulation, a genetic algorithm is implemented to solve the inverse problem and produce an optimized time-temperature profile as an output.

7.4.1 Description of the Genetic Algorithm Implementation

To further speed process development cycle using the simulation, the I2E kinetics simulation was wrapped with a genetic algorithm written in C++ [187] to optimize process conditions for a given impurity distribution. As a form of evolutionary global optimization, the genetic algorithm begins with a population of "individuals" each consisting of a set of decision variables, "genes". In our case, the decision variables are the temperatures as a function of time during the phosphorus diffusion process, and as such are continuous. Each individual is then a time-temperature profile.

The algorithm is initialized with a starting population of 10s of individuals, each individual consisting of a randomly-selected time-temperature profile. For each individual in the population, the algorithm calculates the individual's "fitness", the parameter that will be maximized over generations of the simulation. The carrier lifetime resulting from an I2E simulation using the individual's time-temperature profile is used as the single-objective fitness. After the fitness of each individual in the population is calculated, fit individuals are recombined to form the next generation, keeping portions of the time-temperature profiles from successful individuals intact to pass on to the next generation. Random mutations are injected to promote convergence to global rather than local optima. Over generations, the population thus in theory evolves toward optimal values. After a set number of generations, the simulation is stopped. The operation of the algorithm is shown schematically in Fig. 7-5.

7.4.2 Process Optimization by Inverse Design: Using the Genetic Algorithm

Using the genetic algorithm, it is possible to sweep essentially the entire time-temperature space of phosphorus diffusion in search of optimal profiles. A simple approach is to break the time-temperature profile down into short segments and allow the temperature of each segment to vary. Taking a 55 min diffusion process and leaving 10 min at the end of the process for an exponential "free-cooling" of the wafers to room temper-

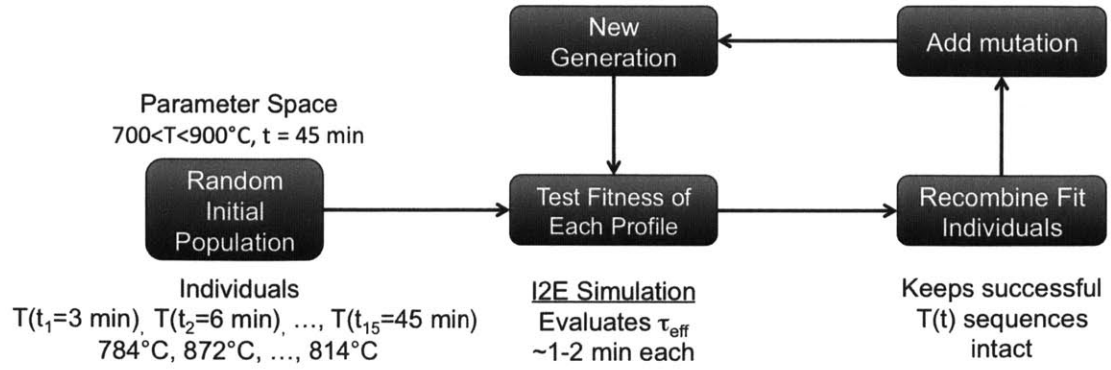

Figure 7-5: A schematic of the application of a genetic algorithm to the optimization of the time-temperature profiles for iron gettering is shown. After initializing the algorithm with a randomly selected population drawn from the entire parameter space, the fitness of each individual (a set of temperatures that together form the time-temperature profile) is tested with the I2E simulation. Based on the resulting lifetimes of the population of profiles run, the fittest individuals are selected for recombination and mutation operations that produce the next generation. Each subsequent generation includes elements of the previous generation, evolving toward a maximum.

ature gives 45 min of high-temperature time. Those 45 min were then divided up into 15 consecutive 3 min intervals and the temperature of each of those intervals allowed to float as a decision variable in the genetic algorithm. The range of temperatures allowed was 700 to 900°C. No constraints were placed to mimic finite ramping times, so the temperature could be 700°C for one interval followed immediately by 900°C the next interval. To investigate gettering for typical multicrystalline material, samples were taken to have 5×10^{14} atoms/cm^3 total Fe (10 ppb) distributed at 35 nm precipitates. An initial population of 60 individual time-temperature profiles was generated randomly, and the algorithm iterated over 50 generations with 90% replacement each generation. Given the random nature of the profiles, including very sharp changes in temperature, integration of the partial differential equations proceeded more slowly, and simulation time increased to roughly 2 min per profile. The 3000-simulation run thus takes a total of roughly 4 days of computational time on a standard PC.

The average lifetime of the population and its standard deviation is shown for each of the 50 generations for two different initializations of the algorithm in Fig. 7-

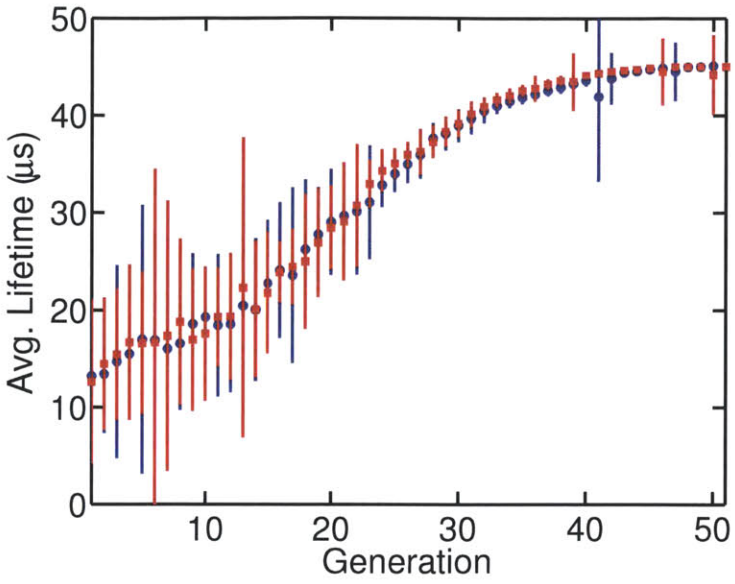

Figure 7-6: The evolution of the average lifetime of the population over the course of the optimization is shown for two instantiations of the genetic algorithm. After a random start, the optimization algorithm converges over 50 generations to the same result in both cases with little final variation in the population, and the lifetime after gettering is greatly improved. Error bars show the standard deviation of the lifetime in the population, and in later generations show the effect of mutations injected into the population. Initial total iron concentration was taken to be 5×10^{14} atoms/cm^3.

6. As can be seen, by the end of 50 generations, the mean resulting lifetime is greatly improved over the initial average and has stabilized at the same value in both instances, with very little variation in the population around it.

Fig. 7-7 shows the 15 temperatures that comprise the best individual profile in each of the 50 generations for the instance of the algorithm shown by the blue dots in Fig. 7-6. The second instance showed similar results. In Fig. 7-7, Temp 1 corresponds to the temperature of the first 3 min interval, Temp 2 the second interval, and so on. One can see that as the algorithm begins there is a high degree of randomness to the best individuals of the first 10 generations as the algorithm covers the 15-dimensional parameter space. Within 20 generations, though, all of the early intervals move toward higher temperatures, while the last two intervals are pushed down to lower temperatures. Generations 40-50 are largely stagnant, with the first 39 min of the profile occurring at the 900°C temperature limit, followed by a step at 750°C and

then a last step at 700°C before the wafers are free-cooled.

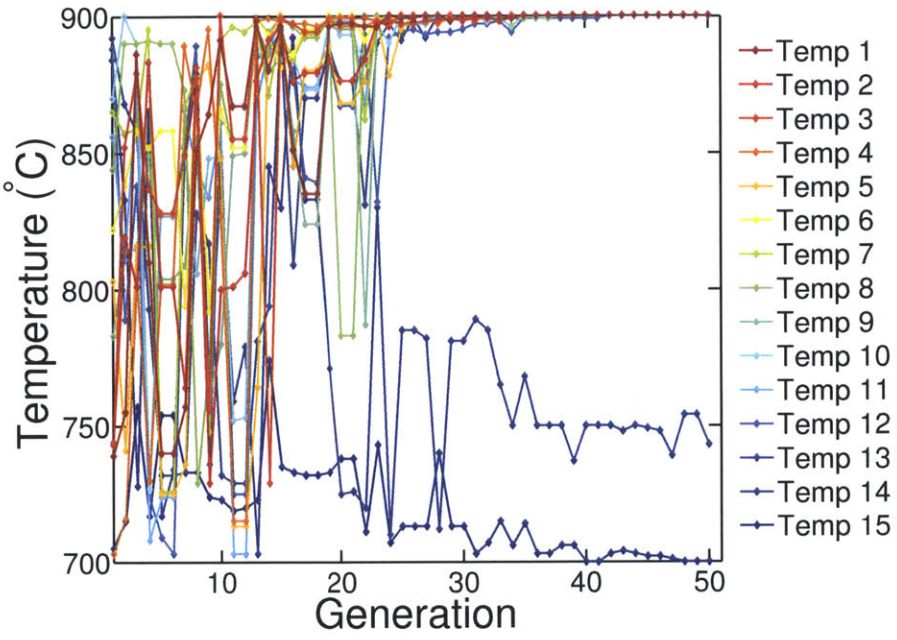

Figure 7-7: The best individual for each generation is plotted, and lines drawn to connect each decision variable of the best individual from generation to generation to track their variation. The temperature variables are listed according to their sequence in the 45 min high-temperature step. During early generations, the best individuals of the population vary widely from generation to generation as the algorithm broadly searches the high-dimensional parameter space. Over time, all of the first 13 temperature intervals converge to the limit at 900°C, while the 14^{th} and 15^{th} steps settle at 750°C and 700°C respectively.

Another visualization of results of the 3000 simulations in a run of the algorithm is achieved by sorting them from poorest lifetime to best without regard to which generation they were from, and plotting every 30th profile, as shown in Fig. 7-8. The lifetime corresponding to each profile is shown in color, with light colors corresponding to poor lifetime, and dark colors corresponding to the best lifetimes. Profiles that remain at 900°C until cooling at the very end of the high-temperature portion of the profile perform best. Once again, no benefit is seen to loading or ramping up slowly from lower temperatures, as shown by simulation and experiment above. The plateau temperature is pushed to the limit at 900°C, maximizing the dissolution of precipitated iron. And the importance of the cooldown is again confirmed, as we see the algorithm produce the best results when the penultimate interval is at 750°C and the

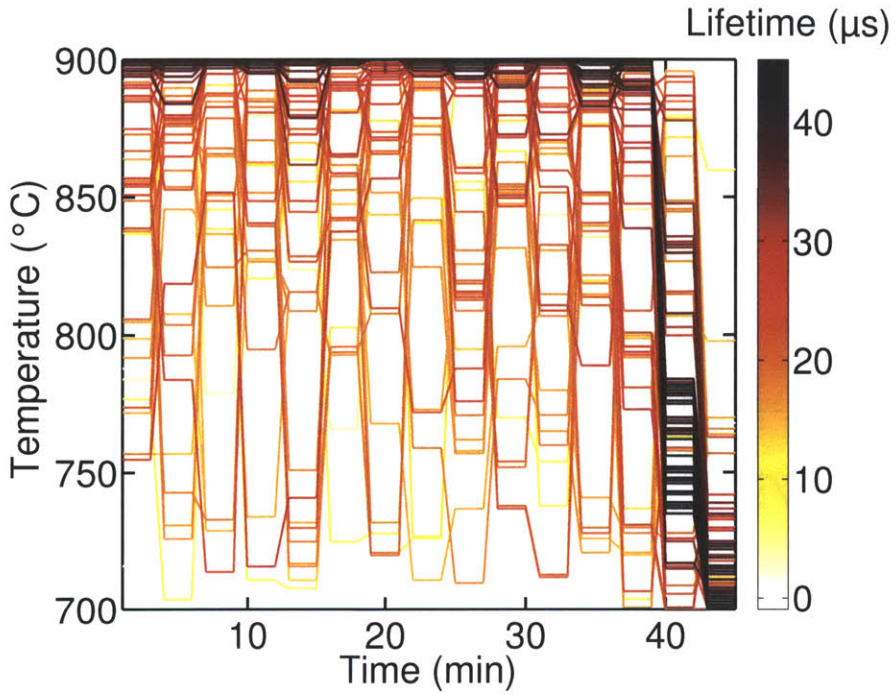

Figure 7-8: A heat map of the phosphorus diffusion parameter space indicates the best profiles (dark lines) are those that run at maximum temperature before cooling steps at the end of the process. The heat map is produced by sorting the 3000 individual simulations run by the genetic algorithm according to final lifetime and plotting every 30th individual. The genetic algorithm has effectively mapped the entirety of the phosphorus diffusion parameter space.

last step is at 700°C. After searching widely across the time-temperature parameter space of phosphorus diffusion with 15 degrees of freedom, the genetic algorithm optimization confirms the design approach discussed above: moving to high temperature to dissolve and getter precipitated iron followed by a controlled, slow cooldown to reduce the interstitial iron concentration.

7.5 Extraction of Design Principles

Fig 7-9 shows the general trends required for improved iron gettering. The high temperature portion of the process is critical to dissolving and removing precipitated iron. Moving to higher gettering temperatures can successfully improve lifetimes in heavily-contaminated materials, as shown in Chapter 6. An unconstrained optimization of the phosphorus diffusion parameter space by the genetic algorithm pushes the

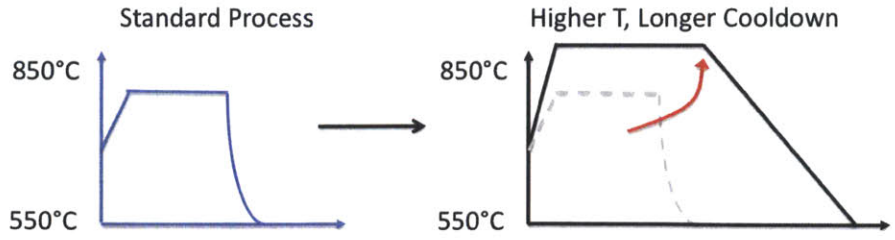

(a) Multicrystalline Si

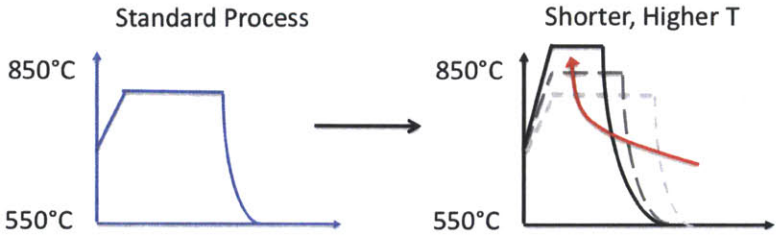

(b) Monocrystalline Si

Figure 7-9: The general trends for improved gettering are shown. (a) In heavily-iron contaminated materials higher gettering temperatures are required to reduce the precipitated iron concentration and improve bulk lifetime. Coupled to a higher gettering temperature must come a controlled cooldown from high temperature so that the interstitial iron concentration is reduced. (b) In mono materials, with less aggressive impurity gettering required, a short, higher-temperature diffusion may be a way of increasing throughput while achieving the same sheet resistance.

gettering temperature to the maximum allowable temperature. But, a higher gettering temperature alone may not improve results, as it can leave a high interstitial iron concentration if the cooldown from the process temperature happens too quickly. Instead, both simulations and experiments show that employing slow-cooling, even for relatively short durations, can dramatically lower the interstitial iron concentration and should be employed when using higher gettering temperatures.

For materials with relatively low total iron concentrations such as typical mono materials, a standard-temperature process can be successful at remediating the iron problem (see *e.g.*, Section 4.2.1). If lifetimes need to be improved further, a slow cooldown or low-temperature anneal may be added to drive enhanced segregation gettering if the lifetime-limiting defect after the standard diffusion remains interstitial

iron. Alternatively, if lifetimes are adequate for the downstream device structure and efficiency targets, a shorter, higher-temperature diffusion may be at least as effective in gettering as a longer, lower-temperature diffusions while producing the same final sheet resistance and decreasing total process time.

7.6 Conclusion

Because the chemical state of iron determines its electronic impact, it is difficult to evaluate *a priori* whether to extend the high-temperature hold time to dissolve additional precipitated iron or favor long cooling steps to force segregation and precipitation of interstitial iron. The tradeoff between higher gettering temperatures and a risk of increased final interstitial iron concentration can be also be difficult to assess. The I2E simulation tool allows us to consider these tradeoffs and predict optimal processes.

First, by enabling a sensitivity analysis, the simulation makes clear which process variables are important to performance. By then performing a one-factor-at-a-time experiment guided by the simulation, the impact of the different pieces of the phosphorus diffusion process on iron gettering become clear, and all trends in the simulation and experimental data agree for both mono- and multicrystalline samples.

By relying on a simple set of equations to describe the most important components of diffusion, segregation, and precipitation of iron during phosphorus diffusion, the impurities-to-efficiency simulation tool allows for large but relatively inexpensive *in silico* optimization experiments. Using the predictive nature of the model in combination with a genetic algorithm, we can explore the wide parameter space of phosphorus diffusion to tailor the diffusion process to the starting material quality to extract the full potential from each wafer.

Chapter 8

Understanding the Limits on Higher-Temperature Processing

Due to the Arrhenius relationship of impurity diffusivity and solubility, higher gettering temperatures should markedly accelerate dissolution of precipitated metals [39]. Confirmation of the improved gettering of precipitated metals at higher temperature was the subject of Chapter 6. However, in the past, mixed results have been seen when performing phosphorus-diffusion gettering at higher temperatures, with material degradation observed by many authors. Peters *et al.* [73] showed that millisecond lifetimes could be maintained in FZ silicon, even with fast ramping rates of roughly 100 K/s using rapid thermal processing, indicating that it is the extrinsic defects in multicrystalline and Czochralski materials that cause degradation. For example, Ballif *et al.* [71] saw lifetime degradation using gettering temperatures above about 875°C in material from the very bottom of a *p*-type cast multi ingot, suggesting that a high dislocation density and dissolving precipitates could be responsible for the increased recombination activity. An order of magnitude higher interstitial iron concentration was seen after diffusion at 950°C compared to a 830°C diffusion. Alternatively, Möller *et al.* [188–190] and Macdonald *et al.* [72] point to the impact of oxygen precipitates, dislocation decoration with metals, and a rise in dislocation density during annealing to account for lower lifetimes after higher-temperature processing.

In this chapter, the trends in lifetime with gettering temperature for the mono- and multicrystalline samples used in this thesis are summarized. The limitations of gettering at higher temperatures are then investigated in the heavily-contaminated multicrystalline silicon materials, by applying detailed microcharacterization to assess potential lifetime-limiting defects, specifically interstitial iron, precipitated iron, and dislocations.

8.1 Lifetime Trends in Different Silicon Materials Processed at High Temperatures

To summarize the effect of process temperature on lifetime for the different materials studied, the resulting area-averaged bulk lifetimes are plotted in Fig. 8-1. The

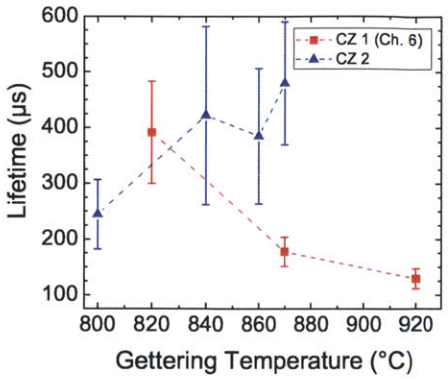

(a) Cz Material Response

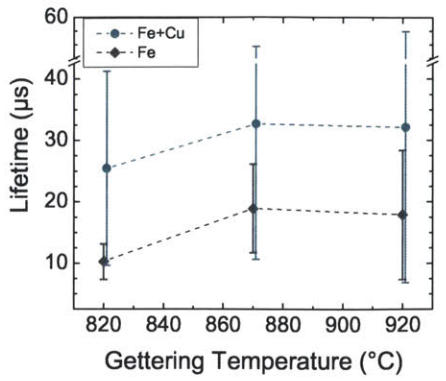

(b) Fe and Fe + Cu mc-Si Response

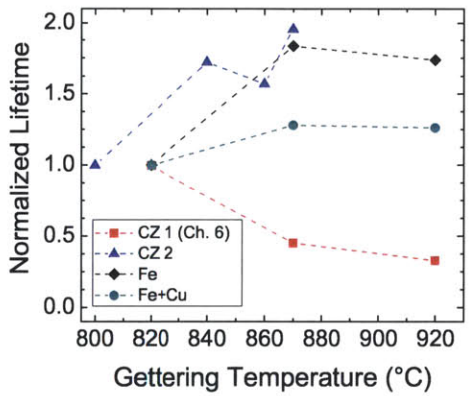

(c) Lifetimes Normalized to Low Temperature Gettering

Figure 8-1: (a) The average lifetime and standard deviation after gettering of two Cz materials is plotted against gettering temperature. Cz 1 is the material used in Ch. 6 which shows strong degradation with higher gettering temperature. It is in stark contrast with Cz 2, a solar Cz material which shows no evidence of degradation, perhaps even some bulk lifetime improvement, with higher-temperature gettering. Error bars show the standard deviation in lifetime across the one Cz 1 sample processed at each temperature, and the standard deviation of average lifetime of the four Cz 2 samples processed at each temperature. (b) The response to higher-temperature gettering indicated by average lifetime and its standard deviation is shown for the Fe and Fe + Cu samples from Ch. 6. The iron-limited material shows significant improvement with higher gettering temperature. Note, however, the different lifetime scale in (b) compared to (a). (c) Arithmetic average lifetimes in the four materials, normalized to the lifetime resulting from an ≤ 820°C P-diffusion step. Fe-contaminated materials improve dramatically with higher temperature diffusions, while higher-purity Cz materials show mixed results.

surfaces of all samples were passivated with iodine ethanol solution for the measurements. The average lifetime and standard deviation as a function of process temperature are shown for two Cz materials in Fig. 8-1(a). Cz 1 is the material used in Ch. 6, measured by μ-PCD, and shows a strong decrease in bulk lifetime with increased gettering temperature. Cz 2, another 1 Ω-cm boron-doped p-type material measured using calibrated-photoluminescence imaging, shows no degradation with temperature, perhaps even showing some improvement though the data is relatively noisy (error bars show the standard deviation of four samples). Given that the injection condition of the μ-PCD measurements are poorly defined, comparison of the absolute lifetime values between the materials is not warranted, but the conflicting trends with temperature are noteworthy. Fig. 8-1(b) shows the lifetime response of the intentionally-contaminated mc-Si materials from Ch. 8 as a function of temperature. Note the decrease of the lifetime scale of roughly one order of magnitude when moving from the Cz materials of Fig. 8-1(a) to Fig. 8-1(b). In response to higher gettering temperatures, the Fe- and Fe + Cu-contaminated materials both respond positively. To emphasize these lifetime trends with temperature, the lifetimes for each material were normalized to the lifetime resulting from the baseline 800 or 820°C gettering process and plotted as a function of process temperature in Fig. 8-1(c).

The Fe-contaminated samples show the greatest benefit from higher gettering temperatures, exhibiting an almost 90% improvement in average lifetime over a baseline gettering process. Much of the improvement in average lifetime is achieved at 870°C, with similar average lifetime but larger variation stemming from a process at 920°C. Lifetime also improves with higher gettering temperatures in the Fe + Cu-contaminated samples, but not to the relative extent achieved in the Fe-only samples. However, as can be seen in Fig 8-1(b), higher absolute lifetimes are measured for a given process temperature in the Fe + Cu-contaminated samples than in Fe-only samples, despite the higher total metal concentrations in the as-grown wafers.

XRF measurements of the precipitated iron distribution in the Fe + Cu samples seen in Fig. 6-6 indicate that the Fe + Cu samples have smaller (or undetectable) iron precipitates. Previously it was shown that dissolution and gettering proceeds

more quickly when smaller precipitates are present [167]. Thus, in materials like the Fe + Cu-contaminated samples, higher-temperature annealing appears to have less of a beneficial effect on lifetime improvement, since lower-temperature $POCl_3$ appears to satisfactorily remediate precipitated iron.

A move toward higher gettering temperatures clearly cannot be recommended for all materials, as, in contrast to the metal-contaminated samples, the higher-purity material Cz 1 shows degradation when process temperature increases. The interstitial oxygen and bulk microdefect concentrations remain to be measured on Cz 1 and Cz 2. The behavior of oxygen precipitates at temperature and their coupling with metal impurities are likely related to the problem, and Murphy *et al.* [191, 192] have begun active investigations in this area. For now, the fundamental material differences that result in opposite trends for Cz 1 and Cz 2 remain the topic of ongoing investigation, and the response of the multicrystalline materials, where bulk iron gettering is of critical importance, is focused on below.

8.2 Spatial Non-Uniformities in mc-Si Lifetime Response: Interactions between Dislocations and Metal Impurities

Upon closer examination of the samples investigated using higher-temperature phosphorus diffusions of Chapter 6, while precipitated iron gettering improves with higher temperatures, some regions of the mc-Si samples degraded with higher-temperature gettering steps leading to large variations in lifetime across the sample (and the large standard deviations seen in Fig. 8-1(b)). For a given mc-Si material, the dislocation density is known to play a crucial role in determining the gettering response with improved gettering often seen in areas of low dislocation density [10, 193], and so areas that getter poorly can be suspected of being rich with structural defects.

8.2.1 Comparing Spatially-Resolved Lifetime, Interstitial Iron, and Dislocation Density

With the knowledge that (1) in addition to iron, heavily-dislocated areas often represent a principal lifetime-limiting defect across much of the area of mc-Si wafers [10] and (2) dislocated areas have been shown to respond poorly to gettering [118, 194], a further investigation of the lifetime limiting defects after higher-temperature phosphorus diffusion was carried out. QSSPC-calibrated photoluminescence imaging was applied to measure spatially-resolved lifetime and interstitial iron concentration [133], followed by Sopori etching and automated etch pit counting to quantify the distribution of defects as a function of gettering temperature. The lifetime, interstitial iron concentration, and etch pit density for the Fe-contaminated samples from Ch. 6 are shown according to their gettering temperature in Fig. 8-2.

In Fig. 8-2(a), one can see that the lifetime after gettering increases with increasing gettering temperature across a significant area of the samples. A close correlation between lifetime and the dislocation density of Fig. 8-2(b) is evident, particularly for the 920°C sample. Enhanced lifetime improvement is seen especially in areas of low dislocation density like the upper and leftmost regions of the samples. However, the heavily dislocated region in the center right of the samples does not improve, and actually degrades somewhat when moving from the 870°C gettering temperature to 920°C.

The interstitial iron concentration measurements of Fig. 8-2(c) provide some insight into the defects underlying the observed lifetime distributions. Focusing first on the 820°C sample, we see that the interstitial iron concentration is distributed rather uniformly across the sample. Even some of the larger grains show a significant interstitial iron concentration, consistent with the relatively poor lifetime performance with respect to the samples gettered at higher temperatures. In contrast, the intragranular regions of the larger grains of the 920°C samples show very low interstitial iron concentrations relatively and consequently much higher lifetimes. On the other hand, turning our attention to the heavily-dislocated region of the sample gettered

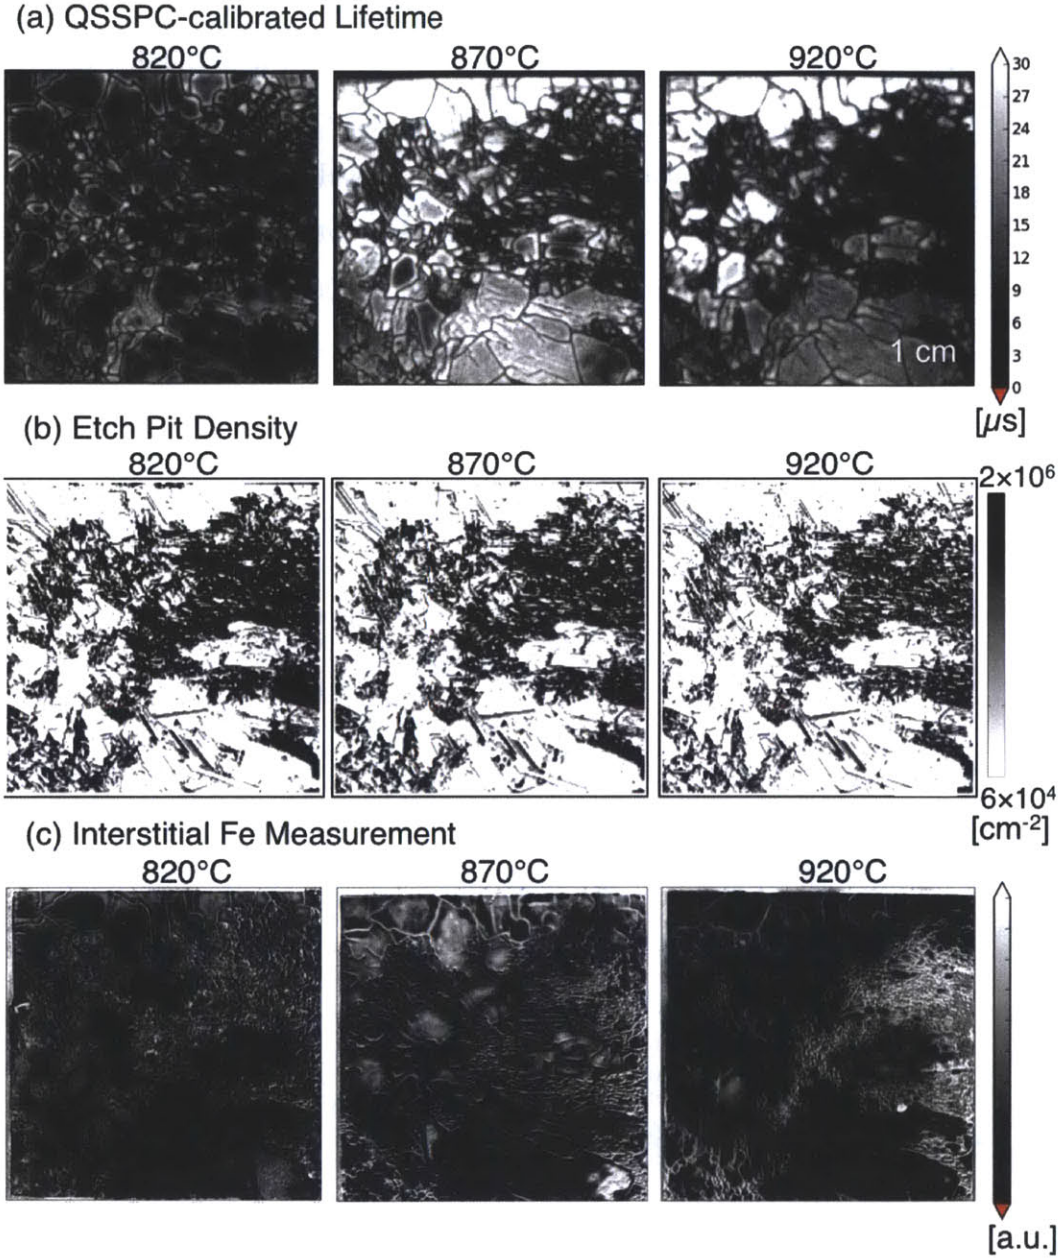

Figure 8-2: (a) QSSPC-calibrated PLI lifetime as a function of sample gettering temperature. Samples are $\approx$ 5×5 cm^2. (b) Quantified etch pit density maps after a 30 s Sopori etch. (c) Corresponding interstitial iron concentration extracted from iron-boron pair dissociation measurements.

at 920°C, it appears that the interstitial iron concentration is especially high there, contributing to the low local lifetimes. It appears that the 920°C phosphorus diffusion process released a significant amount of iron into the surrounding region, perhaps from precipitates that were decorating the dislocations. The high dislocation density and high interstitial iron concentration lead to low lifetimes after gettering in this region. The interstitial iron measurements are displayed in arbitrary units because the injection level achieved in the samples during PLI was near the cross-over point making accurate determination of the C constant difficult (see Section 3.2). Relative comparisons within each sample should be valid, but comparisons between samples must be taken with caution.

8.2.2 Lifetime Histograms as a Function of Processing

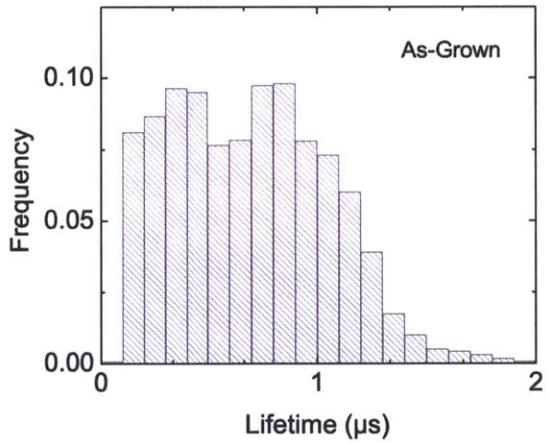

Figure 8-3: The as-grown lifetime frequency distribution is shown for the Fe-contaminated material. Lifetimes are on the order of 1 μs.

To better understand the changes in lifetime as a function of process temperature, histograms of the lifetime distribution are shown for each sister: the sample kept as-grown for reference in Fig. 8-3 and for each gettered sample in Fig. 8-4 according to the gettering process applied. Almost all lifetimes in the as-grown sample are below 2 μs. After gettering at the baseline temperature of 820°C, the lifetime improves dramatically as seen in Fig. 8-4 and is tightly distributed around the mean of 10.3 μs.

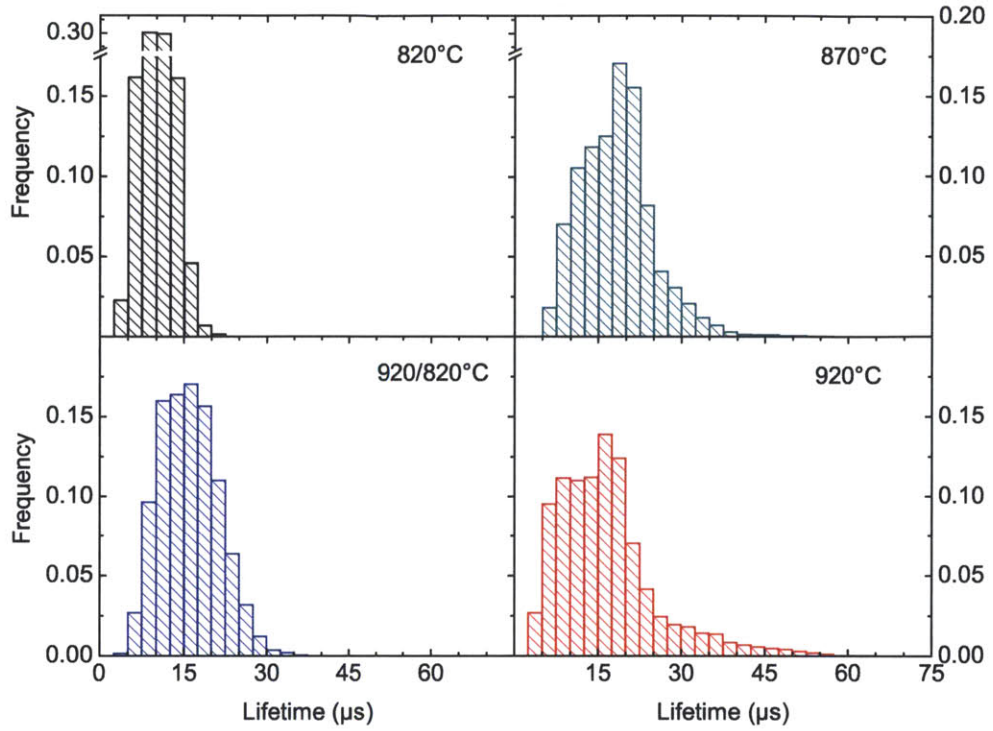

Figure 8-4: Lifetime histograms after gettering. A high-lifetime tail develops in the lifetime distribution by moving to higher-gettering temperatures, but by 920°C this lifetime improvement begins to be offset by degradation in other areas of the sample. The 920°C N_2/820°C PDG improves the mean lifetime while maintaining a tight lifetime distribution.

By moving up the gettering temperature to 870°C, the majority of the lifetime distribution shifts to higher values with a high-lifetime tail developing. From comparison with Fig. 8-2, it is clear that this high-lifetime tail comes from the large-grain areas of low-dislocation density. Moving to the 920°C diffusion produces primarily an increase in variation in lifetime across the sample. The high-lifetime tail is extended with values reaching 75 μs, but a significant fraction of the wafer degrades with respect to the 870°C lifetime distribution, falling below 15 μs. The 920°C N_2 anneal followed by 820°CPDG appears to have achieved its goal of homogenization and improved gettering, with a tight distribution and a higher mean vis-à-vis the 820°C standard diffusion.

8.2.3 Donolato Dislocation Recombination Model as a Framework for Analysis

The lifetime distribution and the lifetime degradation during the 920°C PDG is clearly associated with the presence of dislocations. One way to visualize this local impact of dislocation density is to consider the variation of lifetime as a function of dislocation density.

Donolato [195] developed a dislocation recombination activity model where the dislocation is represented as a cylinder of radius ϵ with a surface recombination velocity for carriers v_s. He determined a relationship between the dislocation-limited lifetime, τ_d, and the dislocation density, ρ_d, that depends only on the parameters ϵ and v_s. These can be collapsed into the non-dimensional dislocation recombination parameter, Γ, by:

$$\Gamma = \frac{2\pi\epsilon v_s}{D} \tag{8.1}$$

where D is the minority carrier diffusivity. The dislocation-limited lifetime, τ_d , is then:

$$\tau_d \cong \frac{1}{2\pi\rho_d D}\left[-\ln(\epsilon\rho_d^{1/2}) - \frac{\ln\pi}{2} - C + \frac{1}{2}\right] + \frac{1}{\Gamma D\rho_d} \tag{8.2}$$

where small approximations have been made to achieve a closed-form solution [195].

A measurement of the effective lifetime can then be broken down into two contributions from the dislocation-limited lifetime, τ_d, and the dislocation-free bulk lifetime, τ_o, according to:

$$\frac{1}{\tau_{eff}} = \frac{1}{\tau_d} + \frac{1}{\tau_o} \tag{8.3}$$

By registering the scanner etch pit density images with the μ-PCD lifetime measurements of the Fe-contaminated samples, a 250 μm resolution point-by-point analysis of the lifetime vs. dislocation density was conducted. The results are shown for the Fe-contaminated samples in Fig. 8-5. A non-linear Levenberg-Marquardt least squares fit (MATLAB, *nlinfit*) was performed to fit the parameters Γ and τ_o for each sample's experimental data according to the Donolato model. The dislocation effective radius, ϵ, is not influential to the parameterization when $\Gamma << 1$ [195] and was

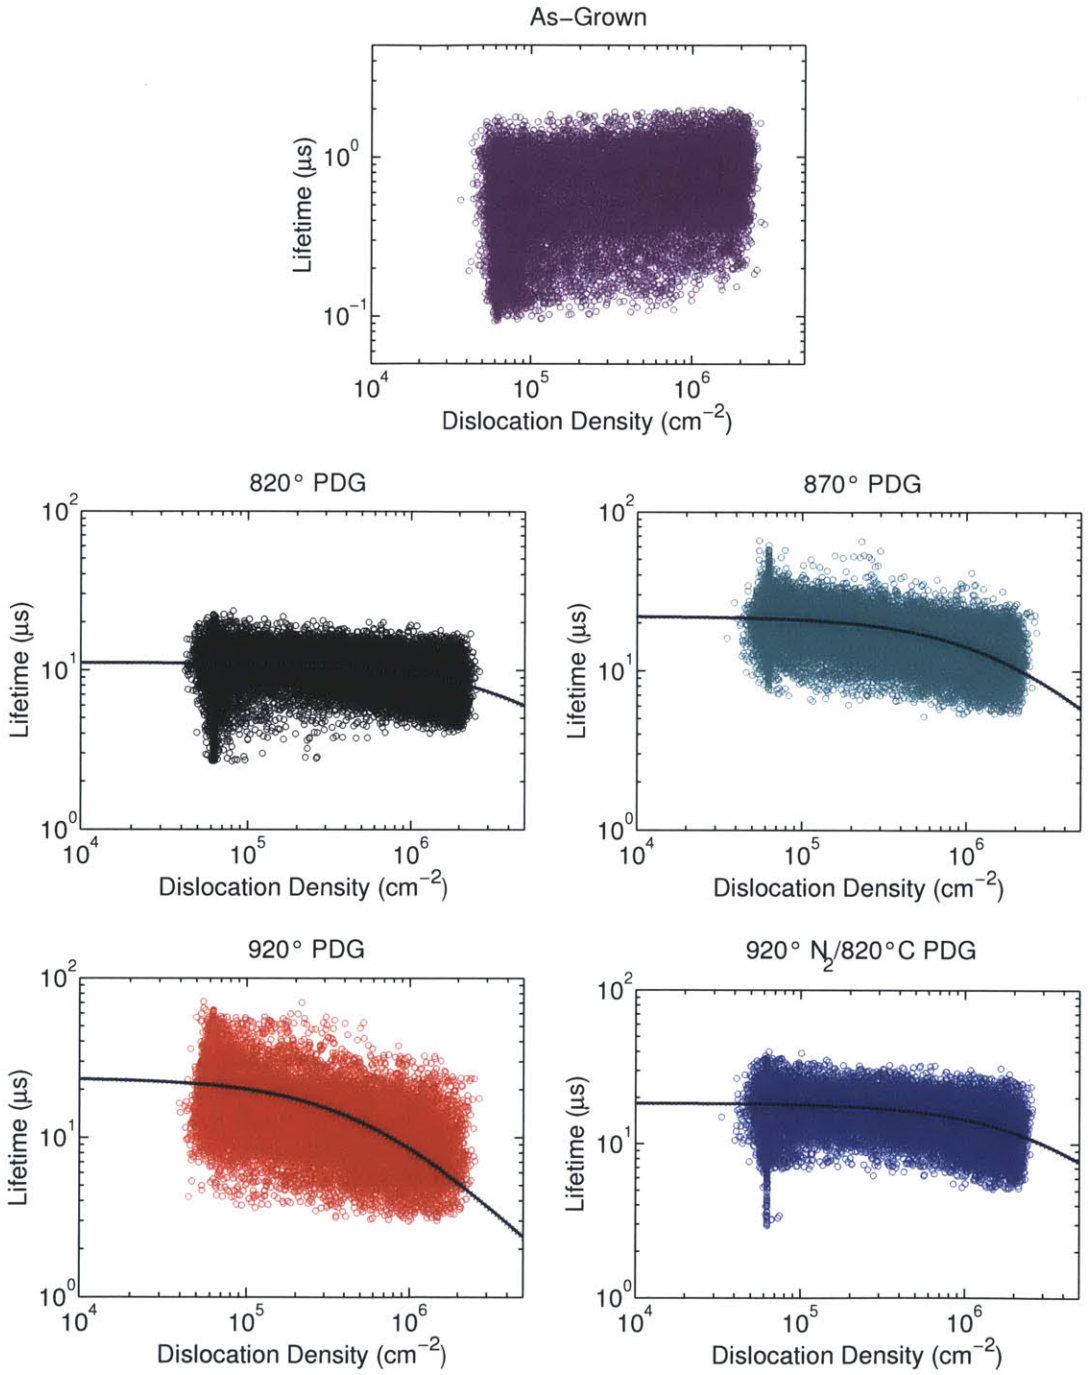

Figure 8-5: The spatially-resolved lifetime is plotted against local dislocation density for the Fe-contaminated samples. The fits resulting from application of Donolato's dislocation recombination model [195] to the data are also shown.

taken to be 50 nm [196]. The results of the fitting, including 95% confidence intervals on the parameters, are shown in Table 8.1.

Table 8.1: The results from fitting the Donolato model to the lifetime vs. dislocation density data are shown according to the different processing applied to the sample.

	As-Grown	820°C	870°C	920°C	920/820°C
Γ (10^{-4})	-106 ± 3	4.6 ± 0.2	7.6 ± 0.2	22.6 ± 0.8	4.6 ± 0.2
τ_o (μs)	0.6 ± .01	11.1 ± 0.1	22.0 ± 0.1	23.7 ± 0.2	18.5 ± 0.1

The as-grown sample shows an increasing lifetime with increasing dislocation density, resulting from the advantages of internal gettering in the low-lifetime sample. The as-grown lifetime distribution (Fig. 8-3) hovers around 1 μs because of high interstitial iron concentration throughout. Near structural defects, where the interstitial iron can diffuse to and precipitate at during the cooldown from crystallization, denuded zones of "high" lifetime form (see the as-grown lifetime map in [166] for comparison with the etch pit density maps in Fig. 8-2). Application of the Donolato model results in a telling, though somewhat nonsensical, negative Γ factor in the as-grown sample as seen in Table 8.1.

The positive bump in dislocation-free lifetime when moving from the 820°C to 870°C gettering is clear in Fig. 8-5, as is the increased variation in lifetime for the 920°C sample. The parameterization with Donolato's model is also helpful. For example, we see that the dislocation-free lifetime, τ_o, increases with increasing gettering temperature, including from 870°C to 920°C, where other measures like the average lifetime decrease. This increase in the dislocation-free lifetime with increasing gettering temperature supports the original idea of increasing the gettering temperature to enhance dissolution and extraction of precipitate iron that was shown in Ch. 6. It is interesting to note that the dislocation-free lifetime of the 920°C N_2/820°C PDG sample does improve upon the dislocation-free lifetime of the 820°C PDG sample, another indication that the 920°C pre-annealing enhances the subsequent phosphorus-diffusion gettering. But the dislocation-free lifetime of the 920°C N_2/820°C PDG sample fails to reach that provided by the higher-temperature PDG.

However, the application of the Donolato model in this case should not be stretched too far. We know from Fig. 8-2 that interstitial iron increases in the heavily-dislocated areas of the 920°C PDG sample, but no accounting for the impact of the increased interstitial iron concentration on the dislocation-free lifetime is made. Instead, the application of the Donolato model yields an increased Γ factor for the dislocations in the 920°C PDG sample to account for the decreased lifetime in dislocation-dense areas. In practice, it is very difficult to separate the impacts of recombination at metal impurities trapped at dislocations or otherwise in their immediate vicinity (perhaps after being released during a high-temperature gettering step like in the 920°C PDG sample) from the recombination activity of the dislocation itself. Efforts to understand the electronic and chemical state of iron as present at dislocations would be most interesting.

8.2.4 Dislocation Decoration with Metal Impurities Correlates with High Recombination Activity

The Donolato model has also been used to show that clusters of dislocations even millimeters apart show widely varying recombination activity [163]. To better understand the differences in recombination activity of localized dislocation clusters, samples from the red zone material studied in Ch. 5 were first characterized by high-resolution laser-beam induced current to measure the internal quantum efficiency (IQE). The IQE gives a measure of local recombination activity and was found to vary strongly between nearby dislocation clusters [146]. By correlating the IQE with the local dislocation density, determined by counting etch pits produced by selective etching [197], the normalized recombination strength for a dislocation cluster, Γ, can be calculated [195, 196]. A high Γ factor means that even a low density of dislocations results in strong recombination.

Two dislocation clusters from the same sample, one cluster with high recombination activity ($\Gamma = 0.012$) and one with low recombination activity ($\Gamma = 0.001$), were investigated at the Hard X-ray Nanoprobe, beamline 26-ID at Argonne's Advanced

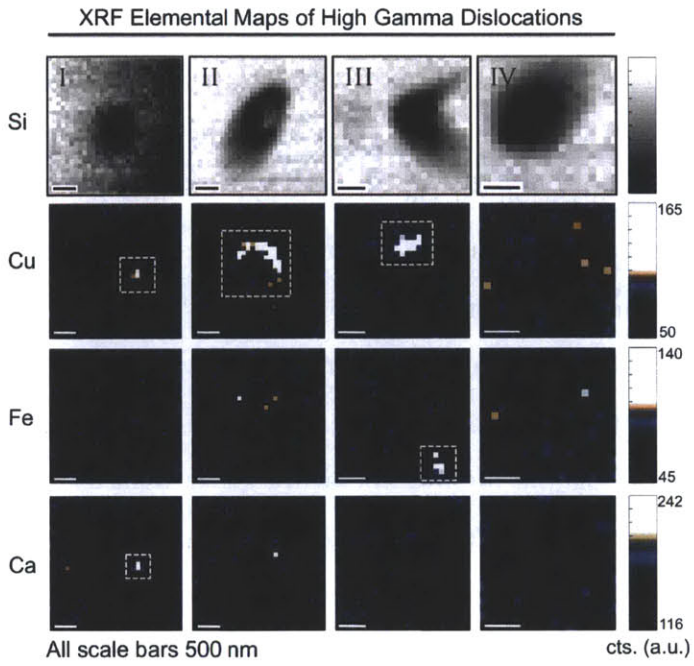

Figure 8-6: Nano-XRF results of highly-recombination active dislocations. The Si channel is used to distinguish the etch pits formed by selective etching at the dislocations. Metal decoration of etch pits is seen within a highly-recombination active dislocation cluster. Cu, Fe, and Ca are observed above detection limits. Reproduced from [146].

Photon Source, using a setup described in Ch. 3. The XRF maps of four dislocations in the high Γ cluster are shown in Fig. 8-6. The etch pit formed at the dislocation can be clearly distinguished in the Si XRF map. The Cu, Fe, and Ca maps show distinct clusters of metals at three of the four etch pits formed. The colormaps of all Cu and Fe elemental XRF maps shown were scaled between one standard deviation (black level) and four standard deviations (white level) above the mean signal over the map area to emphasize clusters where metal concentrations were significantly above the background signal.

Fig. 8-7 shows the XRF results from the low recombination activity dislocation cluster (low Γ). The ten etch pits shown reveal no evidence of metal decoration above detection limits. The XRF results suggest a strong correlation between metal decoration and dislocation recombination activity [146]. Recent nano-XRF measurements of sister wafers that had undergone a complete cell process have confirmed this

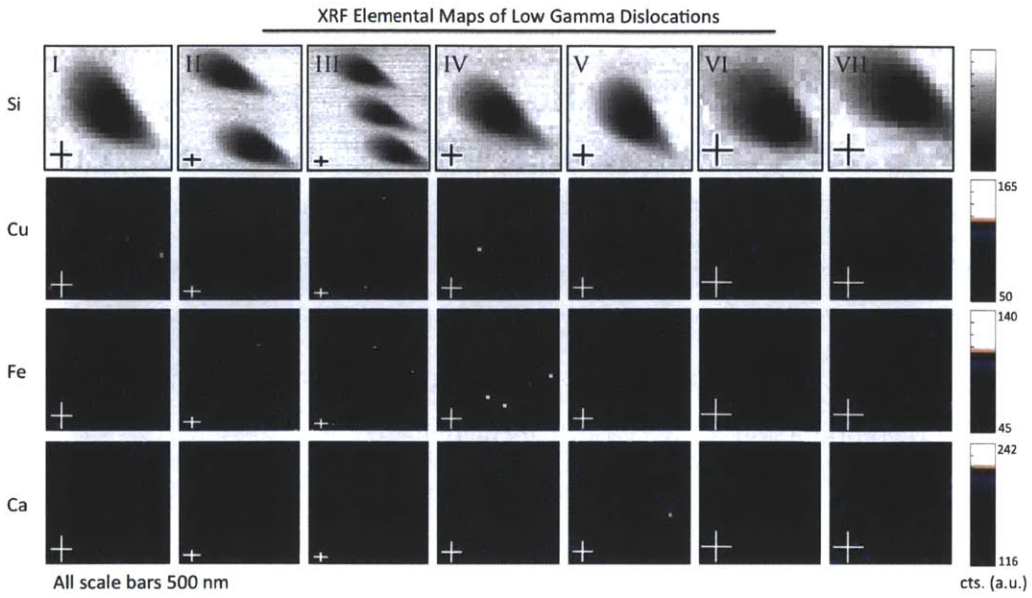

Figure 8-7: No metals above detection limits were observed at dislocations within a cluster showing low recombination activity. Reproduced from [146].

result [198].

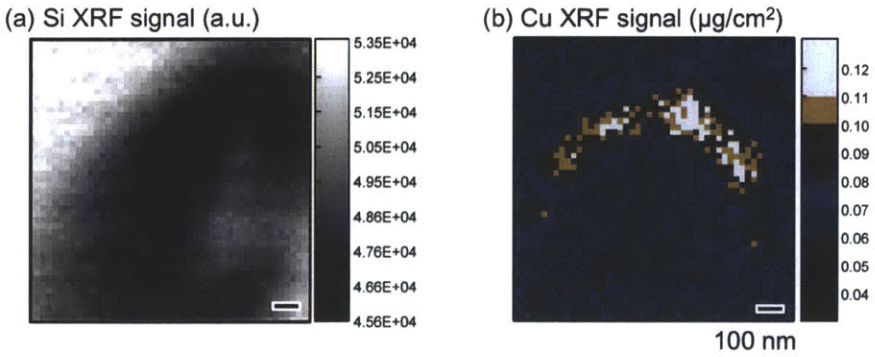

Figure 8-8: Decoration appears quite conformal to the etch pit. Reproduced from [146].

The question remains whether the metal impurities are a significant *cause* of the recombination activity of the dislocations they decorate. For example, perhaps the underlying structural properties of the dislocations determine the electrical activity, while also encouraging metal accumulation. Fig. 8-8 shows a 25 nm step-size map of the Cu decoration of etch pit II in Fig. 8-6 reveals that the decoration closely conforms to the etch pit. The Cu XRF scale bar shows values of loading concen-

trations ($\mu g/cm^2$) estimated from measurements of Si samples and NIST standard reference materials 1832 and 1833 taken at the Advanced Photon Source beamline 2-ID-D. Concentrations should be accurate to within an order of magnitude, due to necessary approximations made in the analysis. The conformal nature of the decoration suggests that the strain field around the dislocation may be an important factor in the decoration. Micro-Raman spectroscopy studies are underway to study this effect. What is clear is that there is a strong correlation between metal decoration and high-recombination activity of dislocations.

8.3 Etch Pit Density Reduction during Phosphorus Diffusion Gettering

It is interesting to note from Fig. 8-2(b) that the etch pit density in the poor lifetime region (center right of a given sample) *decreases* as the gettering temperature increases from 820 to 920°C. Dislocation density reduction at these relatively low temperatures (below 1000°C) has been observed before by Choi *et al.* [199] during annealing of mc-Si blocks in the presence of a gettering layer, but here we use standard wafers. In the following, it is important to consider whether the etch pit density (EPD) corresponds well to the dislocation density. Choi *et al.* [199] used the same Sopori acidic etching chemistry used here and found that the EPD measured for a Secco etch chemistry was the same within experimental error. X-ray tomography measurements to reveal dislocations, similar to those in [200], are planned to confirm that the EPD is a good proxy for dislocation density.

The distributions of the dislocation density in the samples are extracted from the scanner etch pit images of Fig. 8-2(b) and plotted for the as-grown sample in Fig. 8-9 and for the gettered samples in Fig. 8-10. The as-grown distribution is largely bimodal, with a significant fraction of the wafer containing $< 10^5\,cm^{-2}$ and another significant portion containing $> 10^6\,cm^{-2}$. The effects of gettering on this starting dislocation distribution are readily apparent in Fig. 8-10. Gettering at 820°C in-

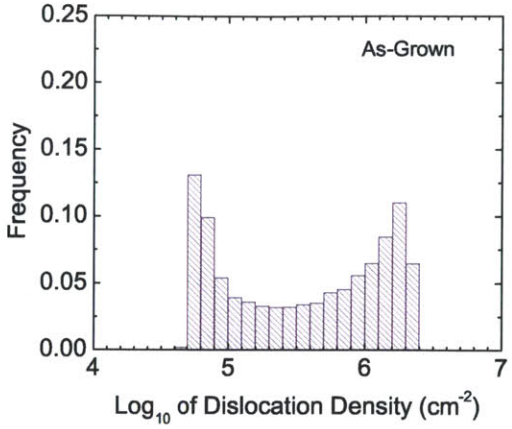

Figure 8-9: The as-grown dislocation frequency distribution is shown on a log scale. The distribution is bimodal, with a large number of pixels at the lower limit of the measurement in the large-grain, largely dislocation-free areas, but a significant fraction of the wafer contains a high density of dislocations ($> 10^6$ cm^{-2})

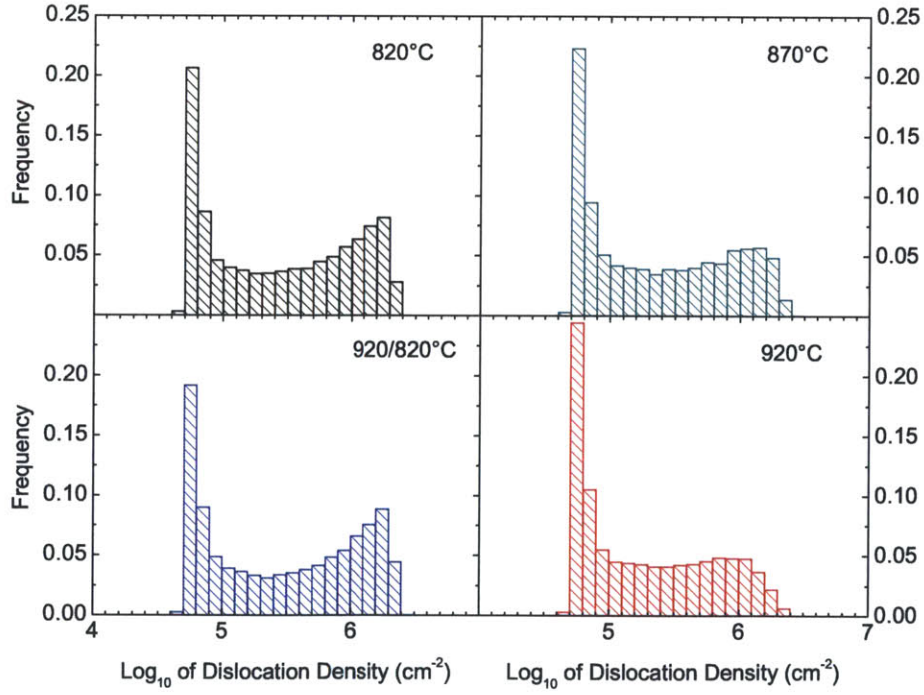

Figure 8-10: Dislocation density histograms shown a log scale. Gettering is seen to reduce the frequency of the highest-density areas, while also increasing the fraction of low-density areas. These trends are monotonic for the gettering processes conducted at 820°C, 870°C, and 920°C. The dislocation density distribution of the 920°C N_2/820°C PDG sample looks most like the 820°C normal PDG sample.

creases the frequency below $10^5\,cm^{-2}$ and decreases the frequency of areas containing dislocation densities above $> 10^6\,cm^{-2}$. Higher-temperature gettering at 870°C and 920°C accentuate these trends. The area fraction in each sample with dislocation density $> 10^6\,cm^{-2}$ is reported in Table 8.2, along with a calculation of the percent reduction in heavily-dislocated area with respect to the as-grown sample. The decrease in high-density ($> 10^6\,cm^{-2}$) area fraction in the 920°C sample is dramatic, with a 64% reduction with respect to the as-grown sample.

Table 8.2: The area fraction of each sample with a dislocation density of more than 10^6 cm^{-2} listed. The reduction of this area fraction as a function of processing is reported. Large reductions in frequency of the most-heavily dislocated areas is seen with higher-temperature P-diffusion steps.

Area with $> 10^6$ cm^{-2} EPD	As-Grown	820°C	870°C	920°C	920/820°C
Area Fraction	0.33	0.25	0.18	0.12	0.28
Reduction from As-Grown	–	24%	44%	64%	16%

To confirm this wafer-level EPD reduction during phosphorus, 20x optical micrographs were taken of a roughly 1.9 cm×1.4 cm region from the center right of the wafers in a region that showed both high and low dislocation-density areas. Fig. 8-11 shows the composite images, automatically stitched together from ≈1.2 mm micrographs, with a 5 mm scale bar. Fig. 8-12 shows portions of the same micrographs, zooming in to the region boxed in Fig. 8-11(v). The scale bar is 1 mm. Dark areas correspond to areas of high EPD. Even a quick inspection of the images reveals a significant reduction in EPD for (ii-v) the diffused samples with respect to (i) the as-grown sample. The most notable reductions occur for the higher-temperature phosphorus diffusions (iii) and (iv). The reduction, comparing the micrographs of (iv) the 920°C sample to (i) the as-grown sample on both the 5 mm and the 1 mm scale, is striking. Importantly, the dislocation distribution in (v) for the sample annealed at 920°C in N_2 prior to an 820°C P-diffusion appears much more similar to (ii) the 820°C PDG sample than (iv) the 920°C PDG, suggesting that it is he temperature of phosphorus diffusion, rather than the annealing temperature that drives

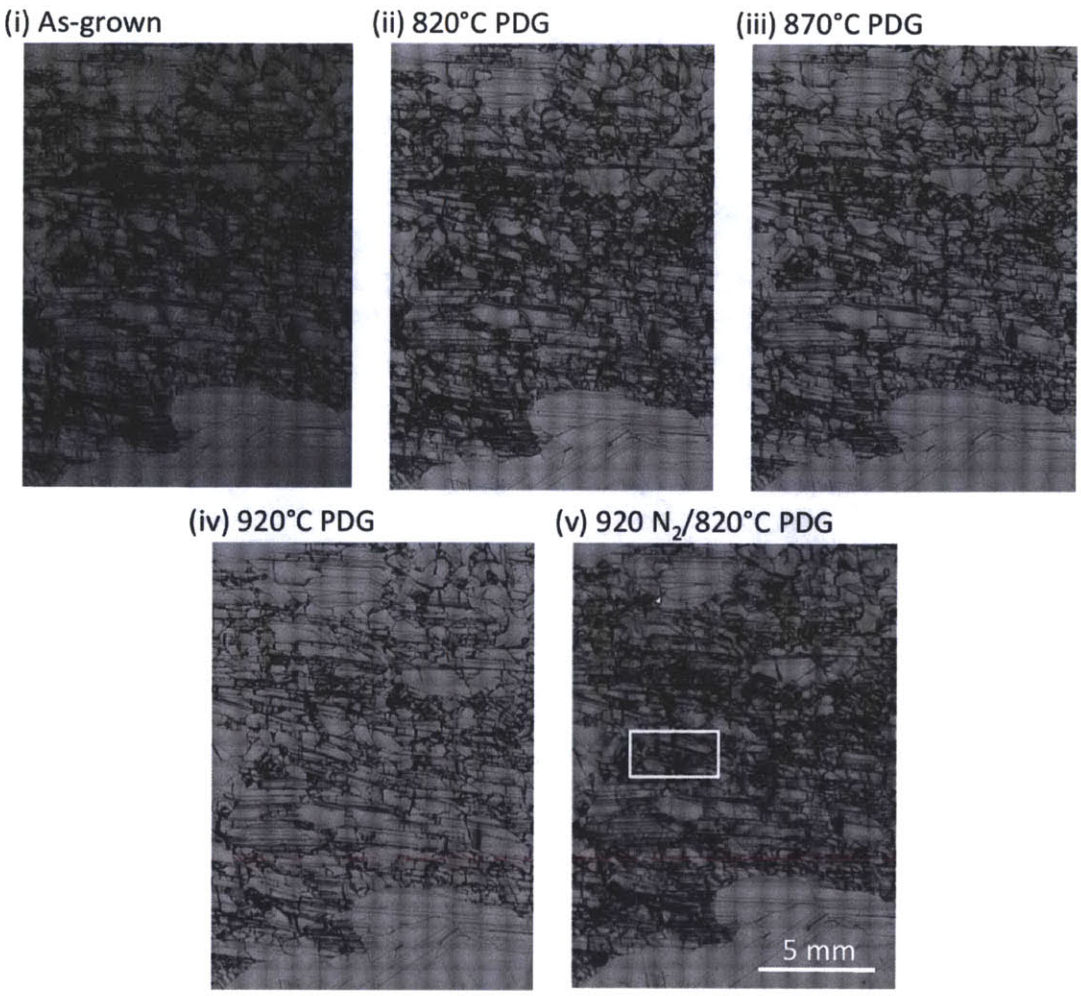

Figure 8-11: Stitched optical micrographs of a 1.9×1.4 cm^2 area of the samples after Sopori etching are shown. The etch pit reduction after processing is pronounced. The reduction appears correlated with the diffusion temperature.

the dislocation reduction.

While the distribution of structural defects, both grain boundaries and dislocations, is remarkably similar in the sister wafers (*cf.* Fig. 8-2), the sister samples are not perfect controls. The samples as ordered from bottom to top in the ingot were as-grown, 820°C, 870°C, 920°C, and 920°C N_2/820°C PDG. The fact that the dislocation density reduction in the 920°C N_2/820°C PDG sample does not continue the trends seen when moving up the ingot with the other four samples, but instead falls in between the 820°C PDG and 920°C PDG response, supports the validity of the results. Further, the Fe + Cu-contaminated samples underwent the same analysis (not

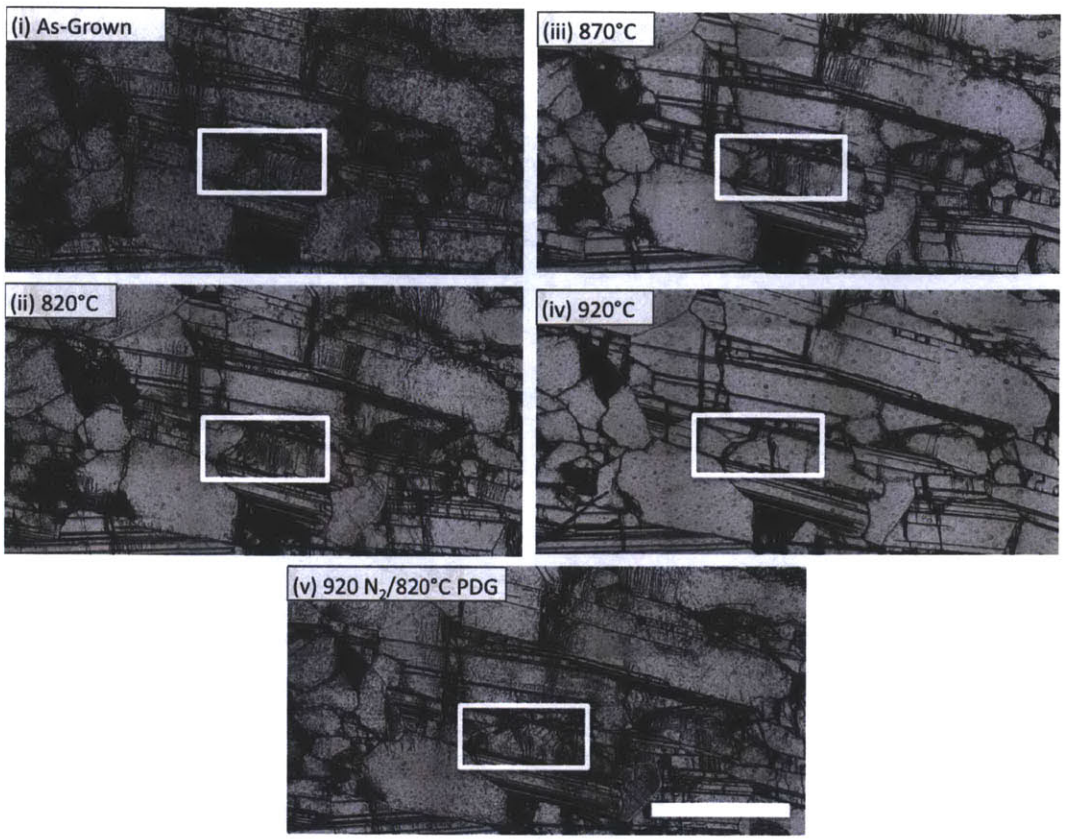

Figure 8-12: Zooming in on the area indicated by the white box in Fig. 8-11 serves to confirm that etch pit reduction is correlated with phosphorus diffusion temperature. The 920°C N_2/820°C PDG appears more similar to the 820°C PDG sample than the 920°C PDG sample. The white boxes indicate the area shown in Fig. 8-13.

shown) and the results were identical in terms of dislocation density reduction as a function of processing.

8.3.1 Tentative Indications of Dislocation Recovery During Phosphorus Diffusion

A final comment on the dislocation redistribution seen during the phosphorus diffusion gettering steps here is the simple magnitude of overall dislocation motion required to produce this reduction and/or annihilation. Fig. 8-13 shows a zoomed in micrograph of the dislocation etch pits in a portion of Fe-contaminated samples. Large differences are seen in the dislocation structures between the samples, beyond the simple reduction of dislocations. Typically it is considered that homologous temperatures

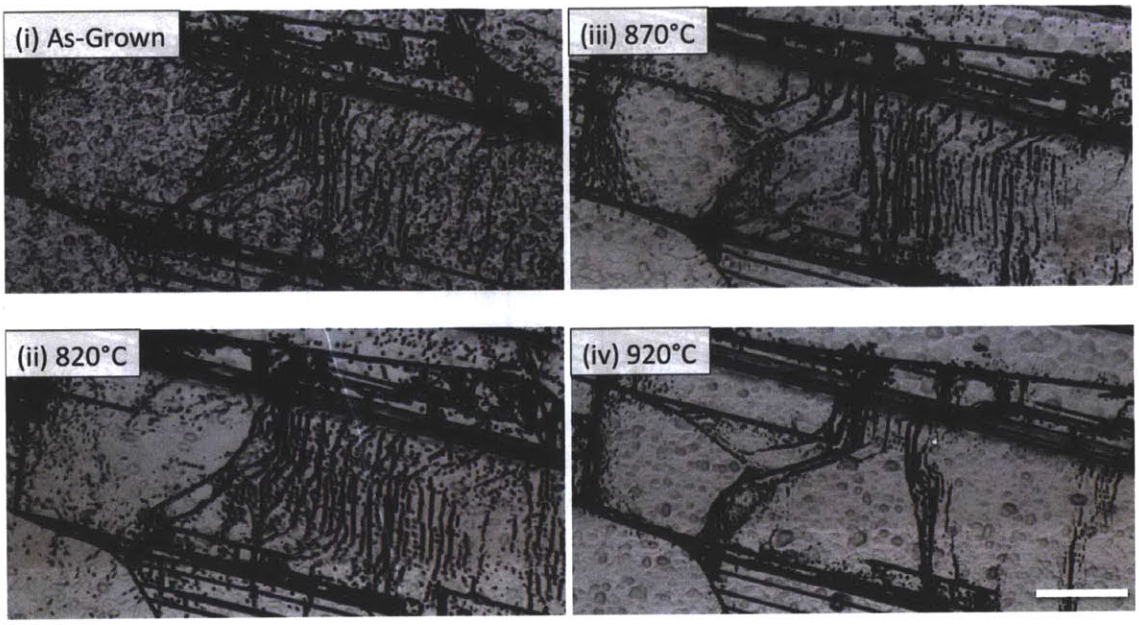

Figure 8-13: Optical micrographs reveal the magnitude of changes in the dislocation structure with the application of a phosphorus diffusion gettering anneal. Higher gettering temperatures lead to large reduction and redistribution. The scale bar is 150 μm.

($T/T_{melting}$) of 0.8 are required for dislocation glide and climb to occur in silicon, and polygonization and dislocation recovery above these temperatures has been observed in multicrystalline silicon [201]. To date, polygonization has not been confirmed here, pending detailed electron backscatter diffraction analysis. However, given that we are operating well below even 1000°C , the significant dislocation redistribution here at phosphorus diffusion temperatures ($0.58 < T/T_m < 0.65$) is surprising. Two speculative mechanisms for the significant changes in the dislocation distribution include (1) motion enhanced by the non-equilibrium injection of Si point defects during phosphorus diffusion and (2) a possible reversal of the typical solute-drag effect of dislocations as depicted in Fig. 8-14. Normally, impurity solute atoms decorate the dislocation core because of the energetically-favorable position in the strain field of the dislocation – the well-known Cottrell atmosphere. Dislocation motion is slowed by the drag imposed by these solute atoms. If during phosphorus diffusion the segregation effect of the gettering layer produces a strong driving force for a net unidirectional flux of

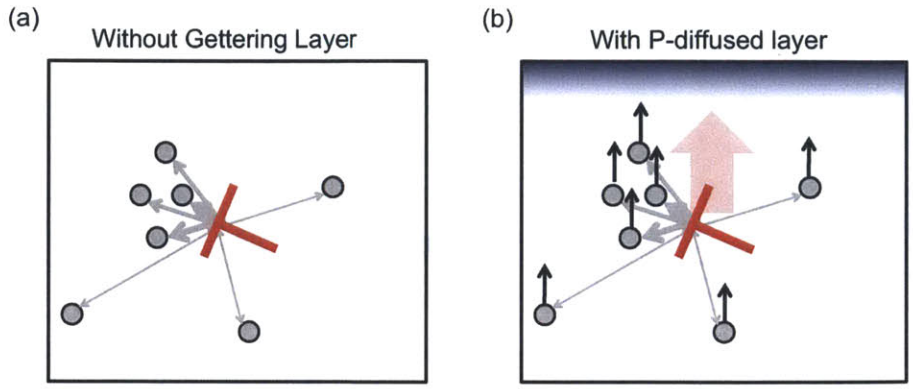

Figure 8-14: (a) Without a gettering layer present, the strain field surrounding a dislocation traps impurity atoms. The presence of impurities along dislocation cores contributes to well-established solute-drag effects. (b) With the introduction of a gettering layer during P-diffusion that provides a strong impurity segregation driving force, a net unidirectional flux of impurities from the bulk to the gettering layer develops. The directional diffusion of these impurities might enhance dislocation motion by the same interaction effects – solute-drag reversed. Adapted from [199].

impurities to the surface layer, then by Newton's third law, one could imaging a net, directional force acting on the dislocation. A solute-drag reversal has been modeled previously using Monte Carlo methods [202], wherein large-magnitude dislocation motion is generated by diffusing solutes. As mentioned above, large dislocation density reduction was observed in thick slabs (5 mm) of mc-Si during phosphorus diffusion by Choi *et al.*, indicating that something other than point defects enhance dislocation motion. The underlying mechanism behind this dislocation motion at the relatively low phosphorus diffusion temperature warrants further investigation.

8.4 Discussion: When to Apply Higher Gettering Temperatures

Overall, the mixed performance in different materials and even within the same material in response to higher-temperature processing demands caution. I suggest a metric in the as-grown material to define whether higher-temperature diffusion could be beneficial: When metal concentrations reach high enough concentrations that denuded zones appear in the vicinity of structural defects with higher lifetimes than

intragranular regions (see *e.g.* [203]). When metals are the principal lifetime-limiting defect in as-grown materials, nucleation and precipitate growth during cooldown from crystallization leads to internal gettering near structural defects, generating higher lifetimes in areas rich with metal impurities and structural defects [133, 204]. During gettering, these precipitates act as semi-infinite sources for metal point defects, poisoning the local area [205]. In contrast, the worst as-grown areas are dominated by high metal point-defect concentrations, due to a paucity of available nucleation sites. This internal gettering effect during cooldown from crystallization leads to the "negative" Γ value at dislocations in the as-grown Fe-contaminated sample above (Fig. 8-5).

During gettering, the extraction of interstitial iron is effective in intragranular regions with low structural defect density. Thus, areas of lower structural defect densities are poor in heavily-contaminated as-grown material but lead to the highest gettered lifetimes: a reversal is seen in the lifetime contrast across structural defects. This appears to be the case with the heavily contaminated Fe-only and Fe + Cu samples, where one observes the lowest as-grown lifetime regions become the highest gettered lifetime regions after a standard diffusion (*cf.* lifetime maps in [166] for full comparison). These regions are largely devoid of dislocations or other structural defects that impede gettering, demonstrated by dislocation etching and imaging seen in Fig. 8-2(b). The lifetime contrast reversal seen here at structural defects does not occur, for example, in a study of metal gettering by Bentzen et al. in higher-quality starting materials [193]. There, the poorest as-grown regions remained the worst areas after gettering. In summary, when lifetime contrast between areas with structural defects and the intragranular bulk reverses after standard gettering, higher-temperature gettering may prove useful because the material is likely metal-limited.

8.5 Conclusions

The potential for lifetime improvement with higher-temperature processing appears to be material dependent – and a specification of the underlying defects driving

the degradation should allow for new tailored defect engineering approaches to be developed that avoid degradation while capturing the benefit of higher temperature processing. To minimize thermal budgets, higher diffusion temperatures should then only be applied when necessary, *i.e.*, to regions of material with high concentrations of Fe-rich precipitates, and must be verified on the cell level.

Fe-limited materials studied here show strong lifetime improvement with higher gettering temperatures, especially in areas of low-dislocation density. Synchrotron-based μ-XRF measurements reveal that gettering of precipitated iron improves at these higher temperatures. However, in heavily dislocated areas, high interstitial iron concentration remains after gettering, limiting lifetime improvement, despite an apparent decrease in the local dislocation density after gettering at higher temperatures. Low-temperature annealing may be effective depending on whether the point defect iron is "trapped" at dislocations. Cz material shows strong degradation with higher gettering temperatures, the mechanism of which is still under investigation. In multicrystalline materials, a reversal in the lifetime constrast between defect-free and defect-rich areas after gettering may be a good indicator of metal-limited material which should respond positively to higher gettering temperatures. A process trade-off can then be determined between improved average lifetime and increased lifetime variance across the wafer to optimize the cell efficiency distribution.

Chapter 9

Conclusion

Today, when a new material with different impurity specifications is introduced into a solar cell manufacturing line, process engineers are faced with high-dimensional, resource-intensive designed experiments to establish a new process. To address this challenge, in this thesis the Impurity-to-Efficiency gettering process simulator is developed to predict final solar cell performance given inputs of material quality and cell processing conditions. The necessary microcharacterization of metal impurities required to populate such models and confirm their predictions is supplied. The main contributions of this thesis include:

- The baselining of the nanoscale response of a contaminated multicrystalline silicon material to a standard phosphorus diffusion gettering. Synchrotron-based X-ray fluorescence is used to identify iron precipitates smaller than 10 nm in radius in materials coming from the bottom, middle, and top of a laboratory-grown ingot. Following I2E model predictions, the experimental data confirm that iron remains in heavily contaminated materials after the high-temperature processing typical of an industrial solar cell process.

- The first large area quantification of the microscopic iron distribution in multicrystalline silicon ingots as a function of position and processing condition using synchrotron-based XRF.

- Establishing that low-temperature annealing, which has been successfully applied to achieve higher solar cell efficiencies in heavily iron-contaminated materials, does not significantly affect the precipitated iron distribution, but strongly reduces the interstitial iron concentration.

- Confirmation that the higher the gettering temperature (up to 920°C), the larger the reduction in precipitated iron in heavily iron-contaminated materials, again by μ-XRF. The large reduction in iron content resulted in an improved minority carrier lifetime.

- Implementation of a genetic algorithm to optimize process conditions for a given impurity distribution. The genetic algorithm has confirmed one-factor-at-a-time

experiments: moving to high temperature to dissolve and getter precipitated iron followed by a controlled, slow cooldown to reduce the interstitial iron concentration results in the highest minority carrier lifetimes.

- Development of design guidelines for phsophorus diffusion tailored to the starting material: the high-temperature portion of the emitter formation process can be designed depending on initial total iron concentration to enhance dissolution of precipitated iron, while the cooldown from the high-temperature process determines the segregation gettering and reprecipitation of dissolved iron. Together the high-temperature step and its cooldown determine the final iron distribution, and the tradeoff between process time and lifetime improvement from iron redistribution can then be made through optimization studies.

- Exploration of the limits of higher-temperature gettering, with evidence that release of iron interstitials from dislocations can lead to degradation in contaminated mc-Si. It remains to be seen whether this released interstitial iron can be extracted from the bulk by application of a low-temperature annealing step following higher-temperature phosphorus diffusion.

- Analysis of etch pit density reduction during phosphorus diffusion gettering, which appears to be correlated with the temperature of phosphorus diffusion, indicating a coupling between either the intrinsic or extrinsic point defect redistribution and dislocation motion.

- The rapid turn-around capability of the I2E model has also allowed for industrial impact. Having confirmed the predictive potential of our model for process improvement, an industrial collaboration guided by the I2E model helped to reduce junction formation time by 5x, while reaching large-area cell efficiencies above 19%. The 5x throughput improvement with improved efficiency greatly reduced tool cost of ownership.

Of course, at least as many questions are raised by this work as are answered. For future work, the interactions between metal impurities and dislocations and their evo-

lution during solar cell processing are surely rich topics for investigation – fundamental materials science questions remain to be answered here. Using EBIC/FIB/TEM studies and possibly adding Atom Probe Tomography to compliment synchrotron-based μ-XRF could be a path forward for investigating the coevolution of metal impurities and structural defects. The competition between internal gettering at structural defects and external gettering to surface layers during processing, especially in response to the rate of cooldown from phosphorus diffusion, may have important implications for final minority carrier lifetime and deserves exploration. The interaction of metal impurities with oxygen precipitates cannot be overlooked either, given the importance of monocrystalline Cz materials in the PV market. This work has focused exclusively on p-type materials, but the impurity kinetics and the impact of second-phase impurity precipitates in n-type silicon materials remain to be explored. Finally, as cost pressure encourages the complete utilization of silicon feedstocks, the question of improving red zone material to produce highly-efficient devices will not disappear. Instead, extended, alternative processing may need to be considered for operation in a parallel stream in the cell manufacturing environment.

More broadly, this thesis has focused on developing an accelerated approach to defect-engineering for renewable energy applications. After the determination of the device-limiting defect, experimental input parameters for their modeling, such as defect distributions and concentrations and defect kinetic reaction rates constants, are collected from the literature or experimentally determined. Here, synchrotron-based μ-X-ray fluorescence was the primary tool to gather the necessary input parameters. The gettering model is reduced to practice and streamlined through the judicious use of engineering assumptions to simplify the governing equations and make practical their numerical solution. Finally, an optimization algorithm is applied to confirm the experimental trends seen and sweep the full parameter space, extracting maximum information from the model. While applied here to silicon solar cells, these concepts for defect modeling are widely applicable to semiconducting materials for thin-films, nanostructure manufacturing, the development of electrochemical storage materials, and general annealing processes and material transformation of energy materials.

Bibliography

[1] W. A. Hermann, "Quantifying global exergy resources," *Energy*, vol. 31, no. 12, pp. 1685–1702, 2006.

[2] Internation Energy Agency, "Trends in Photovoltaic Applications: Survey report of selected IEA countries between 1992 and 2011," *Report IEA-PVPS*, vol. T1-21, pp. 1–48, Sep. 2012.

[3] S. Mehta, "28th Annual Cell and Module Data Collection Results," *PVNews*, vol. 31, no. 5, p. 2, May 2012.

[4] A. A. Istratov, H. Hieslmair, and E. R. Weber, "Iron contamination in silicon technology," *Applied Physics A: Materials Science & Processing*, vol. 70, pp. 489–534, 2000.

[5] G. L. Araujo and A. Marti, "Absolute Limiting Efficiencies for Photovoltaic Energy Conversion," *Solar Energy Materials and Solar Cells*, vol. 33, no. 2, pp. 213–240, 1994.

[6] C. del Cañizo, G. del Coso, and W. C. Sinke, "Crystalline silicon solar module technology: Towards the 1 per watt-peak goal," *Progress in Photovoltaics Research and Applications*, vol. 17, no. 3, pp. 199–209, 2009.

[7] H. J. Queisser and E. E. Haller, "Defects in semiconductors: Some fatal, some vital," *Science*, vol. 281, no. 5379, pp. 945–950, 1998.

[8] J. R. Davis Jr, A. Rohatgi, R. H. Hopkins, P. D. Blais, P. Rai-Choudhury, J. R. McCormick, and H. C. Mollenkopf, "Impurities in silicon solar cells," *IEEE Transactions on Electron Devices*, vol. 27, no. 4, pp. 677–687, 1980.

[9] G. del Coso, C. del Cañizo, and W. C. Sinke, "The impact of silicon feedstock on the PV module cost," *Solar Energy Materials and Solar Cells*, vol. 94, no. 2, pp. 345–349, 2010.

[10] B. Sopori, "Silicon solar-cell processing for minimizing the influence of impurities and defects," *Journal Of Electronic Materials*, vol. 31, no. 10, pp. 972–980, 2002.

[11] J. Hofstetter, J.-F. Lelièvre, D. P. Fenning, M. I. Bertoni, T. Buonassisi, and C. del Cañizo, "Towards the Tailoring of P Diffusion Gettering to As-Grown

Silicon Material Properties," *Solid State Phenomena*, vol. 178-179, pp. 158–165, Aug. 2011.

[12] D. Macdonald, A. Cuevas, and J. Wong-Leung, "Capture cross sections of the acceptor level of iron–boron pairs in p-type silicon by injection-level dependent lifetime measurements," *Journal Of Applied Physics*, vol. 89, no. 12, p. 7932, 2001.

[13] S. Rein and S. W. Glunz, "Electronic properties of interstitial iron and iron-boron pairs determined by means of advanced lifetime spectroscopy," *Journal Of Applied Physics*, vol. 98, no. 11, p. 113711, 2005.

[14] E. R. Weber, "Transition metals in silicon," *Applied Physics A: Materials Science & Processing*, vol. 30, no. 1, pp. 1–22, 1983.

[15] J. D. Murphy and R. J. Falster, "The relaxation behaviour of supersaturated iron in single-crystal silicon at 500 to 750 °C," *Journal Of Applied Physics*, vol. 112, no. 11, p. 113506, 2012.

[16] A. A. Istratov, H. Hieslmair, and E. R. Weber, "Iron and its complexes in silicon," *Applied Physics A: Materials Science & Processing*, vol. 69, no. 1, pp. 13–44, 1999.

[17] T. Buonassisi, A. A. Istratov, M. D. Pickett, M. Heuer, J. P. Kalejs, G. Hahn, M. A. Marcus, B. Lai, Z. Cai, S. M. Heald, T. F. Ciszek, R. F. Clark, D. W. Cunningham, A. M. Gabor, R. Jonczyk, S. Narayanan, E. Sauar, and E. R. Weber, "Chemical natures and distributions of metal impurities in multicrystalline silicon materials," *Progress in Photovoltaics Research and Applications*, vol. 14, no. 6, pp. 513–531, 2006.

[18] E. Olsen and E. J. Øvrelid, "Silicon nitride coating and crucible-effects of using upgraded materials in the casting of multicrystalline silicon ingots," *Progress in Photovoltaics Research and Applications*, vol. 16, no. 2, pp. 93–100, 2008.

[19] T. Buonassisi, A. A. Istratov, M. Heuer, M. A. Marcus, R. Jonczyk, J. Isenberg, B. Lai, Z. Cai, S. Heald, W. Warta, R. Schindler, G. Willeke, and E. R. Weber, "Synchrotron-based investigations of the nature and impact of iron contamination in multicrystalline silicon solar cells," *Journal Of Applied Physics*, vol. 97, p. 074901, 2005.

[20] S. Dubois, O. Palais, P. J. Ribeyron, N. Enjalbert, M. Pasquinelli, and S. Martinuzzi, "Effect of intentional bulk contamination with iron on multicrystalline silicon solar cell properties," *Journal Of Applied Physics*, vol. 102, no. 8, p. 083525, 2007.

[21] G. Coletti, R. Kvande, V. D. Mihailetchi, L. J. Geerligs, L. Arnberg, and E. J. Øvrelid, "Effect of iron in silicon feedstock on p- and n-type multicrystalline silicon solar cells," *Journal Of Applied Physics*, vol. 104, no. 10, p. 104913, 2008.

[22] S. Riepe, I. E. Reis, W. Kwapil, M. A. Falkenberg, J. Schön, H. Behnken, J. Bauer, D. Kreßner-Kiel, W. Seifert, and W. Koch, "Research on efficiency limiting defects and defect engineering in silicon solar cells results of the German research cluster SolarFocus," *Physica Status Solidi (c)*, vol. 8, no. 3, pp. 733–738, 2011.

[23] G. Coletti, P. C. P. Bronsveld, G. Hahn, W. Warta, D. Macdonald, B. Ceccaroli, K. Wambach, N. Le Quang, and J. M. Fernandez, "Impact of Metal Contamination in Silicon Solar Cells," *Advanced Functional Materials*, vol. 21, no. 5, pp. 879–890, 2011.

[24] J. Hofstetter, J. F. Lelièvre, C. del Cañizo, and A. Luque, "Acceptable contamination levels in solar grade silicon: From feedstock to solar cell," *Materials Science & Engineering B*, vol. 159, pp. 299–304, 2009.

[25] G. Coletti, "Sensitivity of state-of-the-art and high efficiency crystalline silicon solar cells to metal impurities," *Progress in Photovoltaics Research and Applications*, early view online, Mar. 2012.

[26] R. Kvande, L. J. Geerligs, G. Coletti, L. Arnberg, M. Di Sabatino, E. J. Øvrelid, and C. C. Swanson, "Distribution of iron in multicrystalline silicon ingots," *Journal Of Applied Physics*, vol. 104, no. 6, pp. 064 905–064 909, 2008.

[27] T. U. Naerland, L. Arnberg, and A. Holt, "Origin of the low carrier lifetime edge zone in multicrystalline PV silicon," *Progress in Photovoltaics Research and Applications*, vol. 17, no. 5, pp. 289–296, 2009.

[28] Y. Boulfrad, G. Stokkan, M. M'Hamdi, E. Øvrelid, and L. Arnberg, "Modeling of Lifetime Distribution in a Multicrystalline Silicon Ingot," *Solid State Phenomena*, vol. 178-179, pp. 507–512, 2011.

[29] S. Riepe, J. Isenberg, C. Ballif, S. W. Glunz, and W. Warta, "Carrier density and lifetime imaging of silicon wafers by infrared lock-in thermography," in *Proceedings of the 17th EUPVSEC*, Munich, Germany, 2001, pp. 1597–1599.

[30] A. Goodrich, T. James, and M. Woodhouse, "Residential, Commercial, and Utility-Scale Photovoltaic (PV) System Prices in the United States: Current Drivers and Cost-Reduction Opportunities," NREL, Tech. Rep. NREL/TP-6A20-53347, Feb. 2012.

[31] D. M. Powell, M. T. Winkler, H. J. Choi, C. B. Simmons, D. Berney Needleman, and T. Buonassisi, "Crystalline silicon photovoltaics: a cost analysis framework for determining technology pathways to reach baseload electricity costs," *Energy & Environmental Science*, vol. 5, no. 3, pp. 5874–5883, 2012.

[32] T. Trupke, K. R. McIntosh, J. W. Weber, W. McMillan, L. Ryves, J. Haunschild, C. Shen, and H. Kampwerth, "Inline Photoluminescence Imaging for Industrial Wafer and Cell Manufacturing," in *Proceedings of the 22nd Workshop on Crystalline Silicon Solar Cells*, Vail, CO, Jul. 2012.

[33] M. D. Pickett and T. Buonassisi, "Iron point defect reduction in multicrystalline silicon solar cells," *Applied Physics Letters*, vol. 92, no. 12, p. 122103, 2008.

[34] T. Buonassisi, A. A. Istratov, M. A. Marcus, B. Lai, Z. Cai, S. M. Heald, and E. R. Weber, "Engineering metal-impurity nanodefects for low-cost solar cells," *Nature Materials*, vol. 4, no. 9, pp. 676–679, 2005.

[35] P. Manshanden and L. J. Geerligs, "Improved phosphorous gettering of multicrystalline silicon," *Solar Energy Materials and Solar Cells*, vol. 90, pp. 998–1012, 2006.

[36] J. Härkonen, V.-P. Lempinen, T. Juvonen, and J. Kylmäluoma, "Recovery of minority carrier lifetime in low-cost multicrystalline silicon," *Solar Energy Materials and Solar Cells*, vol. 73, no. 2, pp. 125–130, 2002.

[37] M. Rinio, A. Yodyunyong, S. Keipert-Colberg, Y. P. B. Mouafi, D. Borchert, and A. Montesdeoca-Santana, "Improvement of multicrystalline silicon solar cells by a low temperature anneal after emitter diffusion," *Progress in Photovoltaics Research and Applications*, vol. 19, no. 2, pp. 165–169, 2011.

[38] L. Carnel, P. Hjems, T. Lu, J. Nyhus, K. Helland, and O. Gjerstad, "Influence of wafer quality on cell performance," in *Proceedings of the 34th IEEE PVSC*, Philadelphia, PA, 2009, pp. 36–38.

[39] P. S. Plekhanov, R. Gafiteanu, U. M. Gösele, and T. Y. Tan, "Modeling of gettering of precipitated impurities from Si for carrier lifetime improvement in solar cell applications," *Journal Of Applied Physics*, vol. 86, no. 5, pp. 2453–2458, 1999.

[40] A. Haarahiltunen, H. Savin, M. Yli-Koski, H. Talvitie, M. I. Asghar, and J. Sinkkonen, "As-grown iron precipitates and gettering in multicrystalline silicon," *Materials Science & Engineering B*, vol. 159, pp. 248–252, 2009.

[41] T. Buonassisi, A. A. Istratov, M. A. Marcus, M. Heuer, M. D. Pickett, B. Lai, Z. Cai, S. M. Heald, and E. R. Weber, "Local measurements of diffusion length and chemical character of metal clusters in multicrystalline silicon," *Solid State Phenomena*, vol. 108, pp. 577–584, 2005.

[42] D. P. Fenning, J. Hofstetter, M. I. Bertoni, S. Hudelson, M. Rinio, J. F. Lelièvre, B. Lai, C. del Cañizo, and T. Buonassisi, "Iron distribution in silicon after solar cell processing: Synchrotron analysis and predictive modeling," *Applied Physics Letters*, vol. 98, no. 16, p. 162103, 2011.

[43] M. Aoki, A. Hara, and A. Ohsawa, "Fundamental properties of intrinsic gettering of iron in a silicon wafer," *Journal Of Applied Physics*, vol. 72, no. 3, pp. 895–898, 1992.

[44] D. Macdonald and L. J. Geerligs, "Recombination activity of interstitial iron and other transition metal point defects in p- and n-type crystalline silicon," *Applied Physics Letters*, vol. 85, no. 18, p. 4061, 2004.

[45] W. Shockley and W. Read, "Statistics of the Recombinations of Holes and Electrons," *Physical Review*, vol. 87, no. 5, pp. 835–842, 1952.

[46] R. N. Hall, "Electron-Hole Recombination in Germanium," *Physical Review*, vol. 87, no. 2, p. 387, 1952.

[47] D. Macdonald and A. Cuevas, "Validity of simplified Shockley-Read-Hall statistics for modeling carrier lifetimes in crystalline," *Physical Review B*, vol. 67, p. 075203, 2003.

[48] C. del Cañizo and A. Luque, "A comprehensive model for the gettering of lifetime-killing impurities in silicon," *Journal Of The Electrochemical Society*, vol. 147, no. 7, pp. 2685–2692, 2000.

[49] J. Schön, H. Habenicht, M. C. Schubert, and W. Warta, "Understanding the distribution of iron in multicrystalline silicon after emitter formation: Theoretical model and experiments," *Journal Of Applied Physics*, vol. 109, no. 6, p. 063717, 2011.

[50] P. S. Plekhanov and T. Tan, "Schottky effect model of electrical activity of metallic precipitates in silicon," *Applied Physics Letters*, vol. 76, no. 25, pp. 3777–3779, 2000.

[51] A. Istratov, H. Hieslmair, O. Vyvenko, E. Weber, and R. Schindler, "Defect recognition and impurity detection techniques in crystalline silicon for solar cells," *Solar Energy Materials and Solar Cells*, vol. 72, no. 1-4, pp. 441–451, 2002.

[52] P. Gundel, M. C. Schubert, F. D. Heinz, W. Kwapil, W. Warta, G. Martinez-Criado, M. Reiche, and E. R. Weber, "Impact of stress on the recombination at metal precipitates in silicon," *Journal Of Applied Physics*, vol. 108, no. 10, p. 103707, 2010.

[53] P. Gundel, M. C. Schubert, W. Kwapil, J. Schön, M. Reiche, H. Savin, M. Yli-Koski, J. A. Sans, G. Martinez-Criado, W. Seifert, W. Warta, and E. R. Weber, "Micro-photoluminescence spectroscopy on metal precipitates in silicon," *Physica Status Solidi RRL*, vol. 3, no. 7-8, pp. 230–232, 2009.

[54] T. Buonassisi, O. F. Vyvenko, A. A. Istratov, E. R. Weber, G. Hahn, D. Sontag, J. P. Rakotoniaina, O. Breitenstein, J. Isenberg, and R. Schindler, "Observation of transition metals at shunt locations in multicrystalline silicon solar cells," *Journal Of Applied Physics*, vol. 95, no. 3, pp. 1556–1561, 2004.

[55] W. Kwapil, P. Gundel, M. C. Schubert, F. D. Heinz, W. Warta, E. R. Weber, A. Goetzberger, and G. Martinez-Criado, "Observation of metal precipitates at prebreakdown sites in multicrystalline silicon solar cells," *Applied Physics Letters*, vol. 95, no. 23, p. 232113, 2009.

[56] M. J. Kerr, "Surface, Emitter and Bulk Recombination in Silicon and Development of Silicon Nitride Passivated Solar Cells," Ph.D. dissertation, Australian National University, May 2002.

[57] A. Richter, F. Werner, A. Cuevas, J. Schmidt, and S. W. Glunz, "Improved Parameterization of Auger Recombination in Silicon," *Energy Procedia*, vol. 27, pp. 88–94, 2012.

[58] M. Kerr and A. Cuevas, "General parameterization of Auger recombination in crystalline silicon," *Journal Of Applied Physics*, vol. 91, no. 4, pp. 2473–2480, 2002.

[59] J. Kang and D. Schroder, "Gettering in silicon," *Journal Of Applied Physics*, vol. 65, no. 8, pp. 2974–2985, 1989.

[60] R. Krain, S. Herlufsen, and J. Schmidt, "Internal gettering of iron in multicrystalline silicon at low temperature," *Applied Physics Letters*, vol. 93, no. 15, p. 152108, 2008.

[61] A. Liu, D. Walter, S. P. Phang, and D. Macdonald, "Investigating Internal Gettering of Iron at Grain Boundaries in Multicrystalline Silicon via Photoluminescence Imaging," *IEEE Journal of Photovoltaics*, vol. 2, no. 4, pp. 479–484, 2012.

[62] D. Gilles, E. R. Weber, and S. Hahn, "Mechanism of internal gettering of interstitial impurities in Czochralski-grown silicon," *Physical Review Letters*, vol. 42, no. 2, pp. 196–199, 1990.

[63] H. Hieslmair, A. Istratov, S. McHugo, and C. Flink, "Gettering of iron by oxygen precipitates," *Applied Physics Letters*, vol. 72, no. 12, pp. 1460–1462, 1998.

[64] R. Falster, D. Gambaro, M. Olmo, M. Cornara, and H. Korb, "The Engineering of Silicon Wafer Material Properties Through Vacancy Concentration Profile Control and the Achievement of Ideal Oxygen Precipitation Behavior," in *Proceedings of the Materials Research Society Symposium*, 1998, pp. 27–35.

[65] C. Flink, H. Feick, S. A. McHugo, W. Seifert, H. Hieslmair, T. Heiser, A. A. Istratov, and E. R. Weber, "Out-diffusion and precipitation of copper in silicon: an electrostatic model," *Physical Review Letters*, vol. 85, no. 23, pp. 4900–4903, 2000.

[66] H. Hieslmair, S. McHugo, and E. R. Weber, "Aluminum backside segregation gettering," in *Proceedings of the 26th IEEE PVSC.* Washington, DC: IEEE, May 1996, pp. 441–444.

[67] D. Abdelbarey, V. Kveder, W. Schröter, and M. Seibt, "Aluminum gettering of iron in silicon as a problem of the ternary phase diagram," *Applied Physics Letters*, vol. 94, no. 6, p. 061912, 2009.

[68] D. Gilles, W. Schroter, and W. Bergholz, "Impact of the electronic structure on the solubility and diffusion of 3d transition elements in silicon," *Physical Review B*, vol. 41, no. 9, pp. 5770–5782, 1990.

[69] M. B. Shabani, T. Yamashita, and E. Morita, "Study of Gettering Mechanisms in Silicon: Competitive Gettering Between Phosphorus Diffusion Gettering and Other Gettering Sites," *Solid State Phenomena*, vol. 131-133, pp. 399–404, 2008.

[70] D. Macdonald, A. Cuevas, A. Kinomura, and Y. Nakano, "Phosphorus gettering in multicrystalline silicon studied by neutron activation analysis," in *Proceedings of the 29th IEEE PVSC*, New Orleans, LA, 2002, pp. 285–288.

[71] C. Ballif, S. Peters, D. Borchert, C. Hässler, J. Isenberg, R. Schindler, W. Warta, and G. Willeke, "Lifetime investigations of degradation effects in processed multicrystalline silicon wafers," in *Proceedings of the 17th EU PVSEC*, Munich, Germany, 2001, p. 1818.

[72] D. Macdonald and A. Cuevas, "The Trade-Off Between Phosphorus Gettering and Thermal Degradation in Multicrystalline Silicon," in *Proceedings of the 16th EU PVSEC*, Glasgow, Scotland, Jun. 2000, pp. 1707–1710.

[73] S. Peters, J. Y. Lee, C. Ballif, D. Borchert, S. W. Glunz, W. Warta, and G. Willeke, "Rapid Thermal Processing: A Comprehensive Classification of Silicon Materials," in *Proceedings of the 29th IEEE PVSC*, New Orleans, LA, May 2002, pp. 214–217.

[74] I. Perichaud, "Gettering of impurities in solar silicon," *Solar Energy Matterials and Solar Cells*, vol. 72, pp. 315–326, 2002.

[75] S. A. McHugo, H. Hieslmair, and E. R. Weber, "Gettering of metallic impurities in photovoltaic silicon," *Applied Physics A: Materials Science & Processing*, vol. 64, no. 2, pp. 127–137, 1997.

[76] A. Bentzen and A. Holt, "Correspondence between sheet resistance and emitter profile of phosphorus diffused emitters from a spray-on dopant," in *Proceedings of the 31st IEEE PVSC*, Lake Buena Vista, FL, 2005, pp. 1153–1156.

[77] Air Products & Chemicals, Inc., *Schumacher $POCl_3$ Application Notes*, 1993.

[78] A. Bentzen, A. Holt, J. S. Christensen, and B. G. Svensson, "High concentration in-diffusion of phosphorus in Si from a spray-on source," *Journal Of Applied Physics*, vol. 99, no. 6, p. 064502, 2006.

[79] K. Tang, E. J. Ovrelid, G. Tranell, and M. Tangstad, "A Thermochemical Database for the Solar Cell Silicon Materials," *Materials Transactions*, vol. 50, no. 8, pp. 1978–1984, Jul. 2009.

[80] F. A. Trumbore, "Solid Solubilities of Impurity Elements in Germanium and Silicon," *Bell System Technical Journal*, vol. 39, pp. 205–233, 1960.

[81] R. Fair and J. Tsai, "A Quantitative Model for the Diffusion of Phosphorus in Silicon and the Emitter Dip Effect," *Journal Of The Electrochemical Society*, vol. 124, no. 7, pp. 1107–1118, 1977.

[82] S. M. Hu, P. Fahey, and R. Dutton, "On models of phosphorus diffusion in silicon," *Journal Of Applied Physics*, vol. 54, no. 12, pp. 6912–6922, 1983.

[83] M. Uematsu, "Simulation of boron, phosphorus, and arsenic diffusion in silicon based on an integrated diffusion model, and the anomalous phosphorus diffusion mechanism," *Journal Of Applied Physics*, vol. 82, no. 5, pp. 2228–2246, 1997.

[84] S. Solmi, A. Parisini, R. Angelucci, A. Armigliato, D. Nobili, and L. Moro, "Dopant and carrier concentration in Si in equilibrium with monoclinic SiP precipitates." *Physical Review B*, vol. 53, no. 12, pp. 7836–7841, 1996.

[85] P. Negrini, D. Nobili, and S. Solmi, "Kinetics of Phosphorus Predeposition in Silicon Using POCl," *Journal Of The Electrochemical Society*, vol. 122, no. 9, pp. 1254–1260, 1975.

[86] R. Chen, H. Wagner, A. Dastgheib-Shirazi, M. Kessler, Z. Zhu, P. P. Altermatt, and S. T. Dunham, "Understanding Coupled Oxide Growth and Phosphorus Diffusion in $POCl_3$ Deposition for Control of Phosphorus Emitter Diffusion," in *Proceedings of the 38th IEEE PVSC*, Austin, TX, Jun. 2012, pp. 213–216.

[87] C. Dubé, B. Tsefrekas, D. Buzby, R. Tavares, W. Zhang, A. Gupta, R. J. Low, W. Skinner, and J. Mullin, "High efficiency selective emitter cells using patterned ion implantation," *Energy Procedia*, vol. 8, pp. 706–711, 2011.

[88] V. Bhosle, C. E. Dube, and N. Bateman, "Value of thermal oxide for ion implant based c-Si solar cells," in *Proceedings of the 38th IEEE PVSC*, Austin, TX, 2012, pp. 2128–2131.

[89] T. Mchedlidze and M. Suezawa, "An iron–phosphorus pair in silicon," *Journal of Physics: Condensed Matter*, vol. 16, pp. 79–84, 2004.

[90] T. Mchedlidze and M. Kittler, "Involvement of iron-phosphorus complexes in iron gettering for n-type silicon," *Physica Status Solidi (a)*, vol. 203, no. 4, pp. 786–791, 2006.

[91] F. Morehead and R. Lever, "Enhanced "tail" diffusion of phosphorus and boron in silicon: Selfinterstitial phenomena," *Applied Physics Letters*, vol. 48, no. 2, pp. 151–153, 1986.

[92] E. Spiecker, M. Seibt, and W. Schröter, "Phosphorous-diffusion gettering in the presence of a nonequilibrium concentration of silicon interstitials: A quantitative model," *Physical Review B*, vol. 55, no. 15, pp. 9577–9583, 1997.

[93] V. Kveder, W. Schröter, A. Sattler, and M. Seibt, "Simulation of Al and phosphorus diffusion gettering in Si," *Materials Science & Engineering B*, vol. 71, pp. 175–181, 2000.

[94] S. M. Myers, M. Seibt, and W. Schröter, "Mechanisms of transition-metal gettering in silicon," *Journal Of Applied Physics*, vol. 88, no. 7, pp. 3795–3819, 2000.

[95] A. Ourmazd and W. Schröter, "Gettering of Metallic Impurities in Silicon," in *Proceedings of the Materials Research Society Symposium.* Cambridge Univ Press, 1985, pp. 25–30.

[96] W. Schröter, V. Kveder, M. Seibt, A. Sattler, and E. Spiecker, "Mechanisms and Computer Modelling of Transition Element Gettering in Silicon," *Solar Energy Materials and Solar Cells*, vol. 72, pp. 299–313, 2002.

[97] M. Seibt, A. Sattler, C. Rudolf, O. Voß, V. Kveder, and W. Schröter, "Gettering in silicon photovoltaics: current state and future perspectives," *Physica Status Solidi (a)*, vol. 203, no. 4, pp. 696–713, 2006.

[98] A. Haarahiltunen, H. Savin, M. Yli-Koski, H. Talvitie, and J. Sinkkonen, "Modeling phosphorus diffusion gettering of iron in single crystal silicon," *Journal Of Applied Physics*, vol. 105, no. 2, p. 023510, 2009.

[99] H. Talvitie, V. Vähänissi, A. Haarahiltunen, M. Yli-Koski, and H. Savin, "Phosphorus and boron diffusion gettering of iron in monocrystalline silicon," *Journal Of Applied Physics*, vol. 109, no. 9, p. 093505, 2011.

[100] S. Dunham, "A Quantitative Model for the Coupled Diffusion of Phosphorus and Points Defects in Silicon," *Journal Of The Electrochemical Society*, vol. 139, no. 9, 1992.

[101] B. C. Trzynadlowski, W. Jiang, R. Chen, and S. T. Dunham, "End-to-end predictive modeling of silicon solar cell performance: From process recipe to device simulation," in *Proceedings of the 37th IEEE PVSC*, Seattle, WA, Jun. 2011, pp. 2952–2955.

[102] B. Tryznadlowski, A. Yazdani, R. Chen, and S. T. Dunham, "Coupled Modeling of Evolution of Impurity/Defect Distribution and Cell Performance," in *Proceedings of the 38th IEEE PVSC*, Austin, TX, Jun. 2012, pp. 217–220.

[103] J. Hofstetter, D. P. Fenning, M. I. Bertoni, J.-F. Lelièvre, C. del Cañizo, and T. Buonassisi, "Impurity-to-efficiency simulator: predictive simulation of silicon solar cell performance based on iron content and distribution," *Progress in Photovoltaics Research and Applications*, vol. 19, no. 4, pp. 487–497, 2011.

[104] H. M. You, U. M. Gösele, and T. Y. Tan, "Simulation of the Transient Indiffusion-Segregation Process of Triply Negatively Charged Ga Vacancies in Gaas and Alas/Gaas Superlattices," *Journal Of Applied Physics*, vol. 74, no. 4, pp. 2461–2470, 1993.

[105] T. Tan, "Mass transport equations unifying descriptions of isothermal diffusion, thermomigration, segregation, and position-dependent diffusivity," *Applied Physics Letters*, vol. 73, no. 18, pp. 2678–2680, 1998.

[106] H. Hieslmair, S. Balasubramanian, A. A. Istratov, and E. R. Weber, "Gettering simulator: physical basis and algorithm," *Semiconductor Science and Technology*, vol. 16, pp. 567–574, 2001.

[107] F. Ham, "Theory of diffusion-limited precipitation," *Journal Of Physics And Chemistry Of Solids*, vol. 6, pp. 335–351, 1958.

[108] J. D. Murphy and R. J. Falster, "Contamination of silicon by iron at temperatures below 800 °C," *Physica Status Solidi (RRL)*, vol. 5, no. 10-11, pp. 370–372, 2011.

[109] D. Macdonald, H. Mäckel, and A. Cuevas, "Effect of gettered iron on recombination in diffused regions of crystalline silicon wafers," *Applied Physics Letters*, vol. 88, no. 9, p. 092105, 2006.

[110] J. Tan, D. Macdonald, N. Bennett, D. Kong, A. Cuevas, and I. Romijn, "Dissolution of metal precipitates in multicrystalline silicon during annealing and the protective effect of phosphorus emitters," *Applied Physics Letters*, vol. 91, no. 4, p. 043505, 2007.

[111] A. A. Istratov, T. Buonassisi, R. J. Mcdonald, A. R. Smith, R. Schindler, J. A. Rand, J. P. Kalejs, and E. R. Weber, "Metal content of multicrystalline silicon for solar cells and its impact on minority carrier diffusion length," *Journal Of Applied Physics*, vol. 94, no. 10, p. 6552, 2003.

[112] D. Clugston and P. Basore, "PC1D version 5: 32-bit solar cell modeling on personal computers," in *Proceedings of the 26th IEEE PVSC*, Anaheim, CA, Oct. 1997, pp. 207–210.

[113] D. M. Powell, D. P. Fenning, J. Hofstetter, J.-F. Lelièvre, C. Del Cañizo, and T. Buonassisi, "TCAD for PV: A fast method foraccurately modelling metal impurityevolution during solar cell processing," *Photovoltaics International*, vol. 15, p. 91, 2012.

[114] A. Haarahiltunen, H. Väinölä, O. Anttila, M. Yli-Koski, and J. Sinkkonen, "Experimental and theoretical study of heterogeneous iron precipitation in silicon," *Journal Of Applied Physics*, vol. 101, no. 4, p. 043507, 2007.

[115] J. Schön, "Modellierung von Prozessschritten zur Umlagerung rekombinationsaktiver Defekte in kristallinem Silizium," Ph.D. dissertation, Universität Konstanz, Apr. 2011.

[116] M. C. Schubert, J. Schön, B. Michl, A. Abdollahinia, and W. Warta, "Modeling Distribution and Impact of Efficiency Limiting Metallic Impurities in Silicon Solar Cells," in *Proceedings of the 38th IEEE PVSC*, Austin, TX, Jun. 2012, pp. 286–291.

[117] D. Franke, "Rise of dislocation density in crystalline silicon wafers during diffusion processing," in *Proceedings of the 3rd World Conference on Photovoltaic Energy Conversion*, Osaka, Japan, May 2003, pp. 1344–1347.

[118] A. Bentzen, B. Svensson, E. Marstein, and A. Holt, "The influence of structural defects on phosphorus diffusion in multicrystalline silicon," *Solar Energy Materials and Solar Cells*, vol. 90, no. 18-19, pp. 3193–3198, 2006.

[119] L. J. Geerligs, Y. Komatsu, I. Roever, K. Wambach, I. Yamaga, and T. Saitoh, "Precipitates and hydrogen passivation at crystal defects in n- and p-type multicrystalline silicon," *Journal Of Applied Physics*, vol. 102, no. 9, p. 093702, 2007.

[120] J. Hofstetter, "Defect engineering strategies for solar grade silicon and their optimization by predictive simulation," Ph.D. dissertation, ETSI Telecomunicación (UPM), 2011.

[121] J. Zhao, A. Wang, and M. A. Green, "Emitter design for high-efficiency silicon solar cells. Part I: Terrestrial cells," *Progress in Photovoltaics Research and Applications*, vol. 1, no. 3, pp. 193–202, 1993.

[122] A. Cuevas and D. Russell, "Co-optimisation of the emitter region and the metal grid of silicon solar cells," *Progress in Photovoltaics Research and Applications*, vol. 8, no. 6, pp. 603–616, 2000.

[123] D. Macdonald, A. Cuevas, A. Kinomura, Y. Nakano, and L. J. Geerligs, "Transition-metal profiles in a multicrystalline silicon ingot," *Journal Of Applied Physics*, vol. 97, no. 3, p. 033523, 2005.

[124] D. M. Powell, M. T. Winkler, A. Goodrich, and T. Buonassisi, "Modeling the Cost and Minimum Sustainable Price of Crystalline Silicon Photovoltaic Manufacturing in the United States," *IEEE Journal of Photovoltaics*, vol. 3, no. 2, pp. 662–668, 2013.

[125] A. Dastgheib-Shirazi, H. Haverkamp, B. Raabe, F. Book, and G. Hahn, "Selective emitter for industrial solar cell production: a wet chemical approach using a single side diffusion process," in *Proceedings of the 23rd European Photovoltaic Solar Energy Conference*, Valencia, Spain, Sep. 2008, pp. 1197–1199.

[126] A.-K. Volk, U. Jaeger, J. Rentsch, and R. Preu, "A Novel Approach for Mask-Free Selective Emitter Technology Combining Laser Doping and Wet Chemical Etch Back," in *Proceedings of the 26th EU PVSEC*, Hamburg, Germany, Sep. 2011, pp. 884–889.

[127] D. Rudolph, K. Peter, A. Meijer, O. Doll, and I. Köhler, "Etch Back Selective Emitter Process with Single POCl 3 Diffusion," in *Proceedings of the 26th EU PVSEC*, Hamburg, Germany, Sep. 2011, pp. 1349–1352.

[128] Impurities-to-Efficiency (I2E) Simulator, online applet, http://pv-i2e.mit.edu, accessed May 2013.

[129] D. V. Lang, "Deep-level transient spectroscopy: A new method to characterize traps in semiconductors," *Journal Of Applied Physics*, vol. 45, p. 3023, 1974.

[130] G. L. Miller, D. V. Lang, and L. C. Kimerling, "Capacitance Transient Spectroscopy," *Annual Review of Materials Science*, vol. 7, pp. 377–448, 1977.

[131] G. Zoth and W. Bergholz, "A fast, preparationfree method to detect iron in silicon," *Journal Of Applied Physics*, vol. 67, no. 11, pp. 6764–6771, 1990.

[132] L. J. Geerligs and D. Macdonald, "Dynamics of light-induced FeB pair dissociation in crystalline silicon," *Applied Physics Letters*, vol. 85, no. 22, p. 5227, 2004.

[133] D. Macdonald, J. Tan, and T. Trupke, "Imaging interstitial iron concentrations in boron-doped crystalline silicon using photoluminescence," *Journal Of Applied Physics*, vol. 103, no. 7, p. 073710, 2008.

[134] D. Macdonald, L. Geerligs, and A. Azzizi, "Iron detection in crystalline silicon by carrier lifetime measurements for arbitrary injection and doping," *Journal Of Applied Physics*, vol. 95, no. 3, pp. 1021–1028, 2004.

[135] M. C. Schubert, H. Habenicht, and W. Warta, "Imaging of Metastable Defects in Silicon," *IEEE Journal of Photovoltaics*, vol. 1, no. 2, pp. 168–173, 2011.

[136] D. Macdonald, T. Roth, P. N. K. Deenapanray, K. Bothe, P. Pohl, and J. Schmidt, "Formation rates of iron-acceptor pairs in crystalline silicon," *Journal Of Applied Physics*, vol. 98, no. 8, p. 083509, 2005.

[137] D. Macdonald, T. Roth, P. N. K. Deenapanray, T. Trupke, and R. A. Bardos, "Doping dependence of the carrier lifetime crossover point upon dissociation of iron-boron pairs in crystalline silicon," *Applied Physics Letters*, vol. 89, no. 14, p. 142107, 2006.

[138] R. A. Sinton and A. Cuevas, "Contactless determination of current–voltage characteristics and minoritycarrier lifetimes in semiconductors from quasi-steady-state photoconductance data," *Applied Physics Letters*, vol. 69, no. 17, pp. 2510–2512, 1996.

[139] T. Trupke, R. A. Bardos, and M. D. Abbott, "Self-consistent calibration of photoluminescence and photoconductance lifetime measurements," *Applied Physics Letters*, vol. 87, no. 18, p. 184102, 2005.

[140] C. Donolato, "Theory of beam induced current characterization of grain boundaries in polycrystalline solar cells," *Journal Of Applied Physics*, vol. 54, p. 1314, 1983.

[141] M. Seibt, M. Griess, A. A. Istratov, H. Hedemann, A. Sattler, and W. Schröter, "Formation and properties of copper silicide precipitates in silicon," *Physica Status Solidi A*, vol. 166, pp. 171–182, 1998.

[142] M. A. Falkenberg and M. Seibt, "Transmission electron microscopy analysis of extended defects in multicrystalline silicon using insitu EBIC/FIB sample preparation," *Physica Status Solidi (c)*, vol. 10, no. 1, pp. 32–35, 2013.

[143] S. A. McHugo, "Release of metal impurities from structural defects in polycrystalline silicon," *Applied Physics Letters*, vol. 71, no. 14, pp. 1984–1986, 1997.

[144] T. Buonassisi, M. A. Marcus, A. A. Istratov, M. Heuer, T. F. Ciszek, B. Lai, Z. Cai, and E. R. Weber, "Analysis of copper-rich precipitates in silicon: Chemical state, gettering, and impact on multicrystalline silicon solar cell material," *Journal Of Applied Physics*, vol. 97, no. 6, p. 063503, 2005.

[145] T. Buonassisi, A. A. Istratov, S. Peters, C. Ballif, J. Isenberg, S. Riepe, W. Warta, R. Schindler, G. Willeke, Z. Cai, B. Lai, and E. R. Weber, "Impact of metal silicide precipitate dissolution during rapid thermal processing of multicrystalline silicon solar cells," *Applied Physics Letters*, vol. 87, no. 12, p. 121918, 2005.

[146] M. I. Bertoni, D. P. Fenning, M. Rinio, V. Rose, M. Holt, J. Maser, and T. Buonassisi, "Nanoprobe X-ray fluorescence characterization of defects in large-area solar cells," *Energy & Environmental Science*, vol. 4, no. 10, p. 4252, 2011.

[147] Z. Cai, B. Lai, W. Yun, I. McNulty, A. Khounsary, J. Maser, P. Ilinski, D. Legnini, E. Trakhtenberg, S. Xu, B. Tieman, G. Wiemerslage, and E. Gluskin, "Performance of a High-Resolution X-ray Microprobe at the Advanced Photon Source," in *AIP CP 521, Synchrotron Radiation Instrumentation: Eleventh US National Conference*, Stanford, CA, Oct. 1999, pp. 31–34.

[148] S. Vogt, J. Maser, and C. Jacobsen, "Data analysis for X-ray fluorescence imaging," *Journal De Physique IV*, vol. 104, pp. 617–622, 2003.

[149] S. Vogt, "MAPS: A set of software tools for analysis and visualization of 3D X-ray fluorescence data sets," *Journal De Physique IV*, vol. 104, pp. 635–638, 2003.

[150] A. Zuschlag and G. Hahn, "Detection limit of XRF measurements at different synchrotron radiation facilities," in *Proceedings of the 26th EU PVSEC*, Hamburg, Germany, 2011, pp. 1547–1549.

[151] P. Y. Dusausoy, J. Protas, R. Wandji, and B. Roques, "Crystal Structure of Iron Disilicide, $FeSi_2\beta$," *Acta Crystallographica Section B-Structural Crystallography And Crystal Chemistry*, vol. B 27, p. 1209, 1971.

[152] T. Buonassisi, A. A. Istratov, M. D. Pickett, M. A. Marcus, T. F. Ciszek, and E. R. Weber, "Metal precipitation at grain boundaries in silicon: Dependence on grain boundary character and dislocation decoration," *Applied Physics Letters*, vol. 89, no. 4, p. 042102, 2006.

[153] J. Chen and T. Sekiguchi, "Carrier Recombination Activity and Structural Properties of Small-Angle Grain Boundaries in Multicrystalline Silicon," *Japanese Journal Of Applied Physics*, vol. 46, no. 10A, pp. 6489–6497, 2007.

[154] M. I. Bertoni, S. Hudelson, B. K. Newman, D. P. Fenning, H. F. W. Dekkers, E. Cornagliotti, A. Zuschlag, G. Micard, G. Hahn, G. Coletti, B. Lai, and T. Buonassisi, "Influence of defect type on hydrogen passivation efficacy in multicrystalline silicon solar cells," *Progress in Photovoltaics Research and Applications*, vol. 19, no. 2, pp. 187–191, 2011.

[155] K. Arafune, E. Ohishi, H. Sai, Y. Terada, Y. Ohshita, and M. Yamaguchi, "Study on Iron Distribution and Electrical Activities at Grain Boundaries in Polycrystalline Silicon Substrate for Solar Cells," *Japanese Journal Of Applied Physics*, vol. 45, no. 8A, pp. 6153–6156, 2006.

[156] I. Takahashi, N. Usami, H. Mizuseki, Y. Kawazoe, G. Stokkan, and K. Nakajima, "Impact of type of crystal defects in multicrystalline Si on electrical properties and interaction with impurities," *Journal Of Applied Physics*, vol. 109, no. 3, p. 033504, 2011.

[157] A. Zuschlag, S. Ohl, J. Bernhard, H. Morhenn, J. Ebser, J. Junge, S. Seren, and G. Hahn, "XRF Investigations of the Influence of Solar Cell Processing Steps on Iron and Copper Precipitates in Multicrystalline Silicon," in *Proceedings of the 35th IEEE PVSC*, Honolulu, HI, Jul. 2010, pp. 347–351.

[158] J.-F. Lelièvre, J. Hofstetter, A. Peral, I. Hoces, F. Recart, and C. del Cañizo, "Dissolution and gettering of iron during contact co-firing," *Energy Procedia*, vol. 8, pp. 257–262, 2011.

[159] D. P. Fenning, J. Hofstetter, M. I. Bertoni, G. Coletti, B. Lai, C. del Cañizo, and T. Buonassisi, "Precipitated iron: A limit on gettering efficacy in multicrystalline silicon," *Journal Of Applied Physics*, vol. 113, no. 4, p. 044251, 2013.

[160] G. Coletti, P. C. P. Bronsveld, R. Kvande, L. J. Geerligs, and C. C. Swanson, "Effect of metals in silicon feedstock on multicrystalline," presented at *Arriving at Well-Founded Specifications of Solar Grade Silicon*, Amsterdam, The Netherlands, Nov. 2008.

[161] B. Michl, J. Schön, F. Schindler, W. Warta, and M. C. Schubert, "Reducing the Impact of Iron on Carrier Lifetimes In Multicrystalline Silicon Wafers by Optimized Diffusion Temperature Profiles," in *Proceedings of the 27th EUPVSEC*, Franfurt, Germany, Sep. 2012, pp. 709–713.

[162] J. Schön, A. Haarahiltunen, H. Savin, D. P. Fenning, T. Buonassisi, W. Warta, and M. C. Schubert, "Analyses of the Evolution of Iron-Silicide Precipitates in Multicrystalline Silicon During Solar Cell Processing," *IEEE Journal of Photovoltaics*, vol. 3, no. 1, pp. 131–137, 2012.

[163] M. Rinio, A. Yodyungyong, S. Keipert-Colberg, D. Borchert, and A. Montesdeoca-Santana, "Recombination in ingot cast silicon solar cells," *Physica Status Solidi (a)*, vol. 208, no. 4, pp. 760–768, 2011.

[164] G. Coletti, private communication (2012).

[165] V. Osinniy, A. Nylandsted Larsen, E. Hvidsten Dahl, E. Enebakk, A.-K. Søiland, R. Tronstad, and Y. Safir, "Gettering improvements of minority-carrier lifetimes in solar grade silicon," *Solar Energy Materials and Solar Cells*, vol. 101, pp. 123–130, 2012.

[166] D. P. Fenning, A. Zuschlag, M. I. Bertoni, B. Lai, G. Hahn, and T. Buonassisi, "Improved Iron Gettering of Contaminated Multicrystalline Silicon by High-Temperature Phosphorus Diffusion," *Journal Of Applied Physics*, (accepted).

[167] J. Hofstetter, D. P. Fenning, J.-F. Lelièvre, C. Del Cañizo, and T. Buonassisi, "Engineering metal precipitate size distributions to enhance gettering in multicrystalline silicon," *Physica Status Solidi (a)*, vol. 209, no. 10, pp. 1861–1865, 2012.

[168] T. Mankad, R. Sinton, A. Ristow, N. Akil, and E. Picard, "Inline Bulk Lifetime Prediction using QSSPC," in *Proceedings of the 21st Workshop on Crystalline Silicon Solar Cells*, Breckenridge, CO, Aug. 2011.

[169] S. Hudelson, B. Newman, S. Bernardis, D. Fenning, M. Bertoni, M. Marcus, S. Fakra, B. Lai, and T. Buonassisi, "Retrograde Melting and Internal Liquid Gettering in Silicon," *Advanced Materials*, vol. 22, no. 35, pp. 3948–3953, 2010.

[170] X. W. Lin, J. Washburn, Z. Liliental-Weber, and H. Bernas, "Coarsening and Phase-Transition of $FeSi_2$ Precipitates in Si," *Journal Of Applied Physics*, vol. 75, no. 9, pp. 4686–4694, 1994.

[171] J. Hofstetter, J. F. Lelièvre, C. del Cañizo, and A. Luque, "Study of internal versus external gettering of iron during slow cooling processes for silicon solar cell fabrication," *Solid State Phenomena*, vol. 156-158, pp. 387–393, 2010.

[172] M. Rinio, A. Yodyunyong, M. Pirker, C. Zhang, D. Guenther, P. Botchak, S. Keipert, D. Borchert, M. Heuer, and A. Montesdeoca-Santana, "New Results using a Low Temperature Anneal In Proccesing of Multicrystalline Solar Cells," in *Proceedings of the 24th EU PVSEC*, Hamburg, Germany, Sep. 2009, pp. 1816–1819.

[173] J. Hofstetter, J.-F. Lelièvre, D. P. Fenning, M. I. Bertoni, T. Buonassisi, A. Luque, and C. del Cañizo, "Enhanced iron gettering by short, optimized lowtemperature annealing after phosphorus emitter diffusion for industrial silicon solar cell processing," *Physica Status Solidi (c)*, vol. 8, no. 3, pp. 759–762, 2011.

[174] D. Macdonald, S. P. Phang, F. E. Rougieux, S. Y. Lim, D. Paterson, D. L. Howard, M. D. de Jonge, and C. G. Ryan, "Iron-rich particles in heavily contaminated multicrystalline silicon wafers and their response to phosphorus gettering," *Semiconductor Science and Technology*, vol. 27, p. 125016, 2012.

[175] K. L. Pollock, J. Junge, and G. Hahn, "Detailed Investigation of Surface Passivation Methods for Lifetime Measurements on P-Type Silicon Wafers," *IEEE Journal of Photovoltaics*, vol. 2, no. 1, pp. 1–6, 2012.

[176] E. L. Lehmann and J. M. D'Abrera, *Nonparametrics: statistical methods based on ranks.* New York: Springer, 2006.

[177] S. Holm, "A simple sequentially rejective multiple test procedure," *Scandinavian Journal of Statistics*, vol. 6, no. 2, pp. 65–70, 1979.

[178] A. A. Istratov, T. Buonassisi, M. D. Pickett, M. Heuer, and E. R. Weber, "Control of metal impurities in "dirty" Multicrystalline Silicon for Solar Cells," *Materials Science & Engineering B*, vol. 134, no. 2-3, pp. 282–286, 2006.

[179] J. Isenberg, J. Dicker, and W. Warta, "Averaging of laterally inhomogeneous lifetimes for one-dimensional modeling of solar cells," *Journal Of Applied Physics*, vol. 94, no. 6, pp. 4122–4130, 2003.

[180] A. Bentzen, G. Schubert, J. S. Christensen, B. G. Svensson, and A. Holt, "Influence of temperature during phosphorus emitter diffusion from a spray-on source in multicrystalline silicon solar cell processing," *Progress in Photovoltaics Research and Applications*, vol. 15, no. 4, pp. 281–289, 2006.

[181] M. Heuer, T. Buonassisi, M. A. Marcus, A. A. Istratov, M. D. Pickett, T. Shibata, and E. R. Weber, "Complex intermetallic phase in multicrystalline silicon doped with transition metals," *Physical Review B*, vol. 73, no. 23, p. 5, 2006.

[182] T. Buonassisi, M. Heuer, A. A. Istratov, M. D. Pickett, M. A. Marcus, B. Lai, Z. Cai, S. M. Heald, and E. R. Weber, "Transition metal co-precipitation mechanisms in silicon," *Acta Materialia*, vol. 55, no. 18, pp. 6119–6126, 2007.

[183] H.-W. Guo and S. T. Dunham, "Accurate modeling of copper precipitation kinetics including Fermi level dependence," *Applied Physics Letters*, vol. 89, no. 18, p. 182106, 2006.

[184] D. P. Fenning, M. I. Bertoni, and T. Buonassisi, "Advancements in Synchrotron-Based Characterization of Metals in Silicon," presented at the *5th International Workshop of Crystalline Silicon Solar Cells*, Boston, MA, Dec. 2011 (unpublished).

[185] B. Hoex, S. B. S. Heil, E. Langereis, M. C. M. van de Sanden, and W. M. M. Kessels, "Ultralow surface recombination of c-Si substrates passivated by plasma-assisted atomic layer deposited Al_2O_3," *Applied Physics Letters*, vol. 89, no. 4, p. 042112, 2006.

[186] B. Hoex, "Functional Thin Films for High-Efficiency Solar Cells," Ph.D. dissertation, Technische Universiteit Eindhoven, 2008.

[187] K. Sastry, "Single and Multiobjective Genetic Algorithm Toolbox for Matlab in C++," Illinois Genetic Algorithms Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, Tech. Rep. 2007017, Jun. 2007.

[188] H. J. Möller, T. Kaden, S. Scholz, and S. Würzner, "Improving solar grade silicon by controlling extended defect generation and foreign atom defect interactions," *Applied Physics A*, vol. 96, no. 1, pp. 207–220, 2009.

[189] H. J. Möller, C. Funke, M. Rinio, and S. Scholz, "Multicrystalline silicon for solar cells," *Thin Solid Films*, vol. 487, no. 1-2, pp. 179–187, 2005.

[190] H. Möller, C. Funke, A. Lawerenz, S. Riedel, and M. Werner, "Oxygen and lattice distortions in multicrystalline silicon," *Solar Energy Materials and Solar Cells*, vol. 72, no. 1, pp. 403–416, 2002.

[191] J. D. Murphy, K. Bothe, R. Krain, M. Olmo, V. V. Voronkov, and R. J. Falster, "Recombination at Oxide Precipitates in Silicon," *Solid State Phenomena*, vol. 178-179, pp. 205–210, 2011.

[192] J. D. Murphy, K. Bothe, M. Olmo, V. V. Voronkov, and R. J. Falster, "The effect of oxide precipitates on minority carrier lifetime in p-type silicon," *Journal Of Applied Physics*, vol. 110, no. 5, p. 053713, 2011.

[193] A. Bentzen and A. Holt, "Overview of phosphorus diffusion and gettering in multicrystalline silicon," *Materials Science & Engineering B*, vol. 159, pp. 228–234, 2009.

[194] O. Schultz, S. W. Glunz, S. Riepe, and G. P. Willeke, "Gettering of Multicrystalline Silicon for High-Efficiency Solar Cells," in *Proceedings of the 21st EUPVSEC*, Dresden, Germany, 2006, pp. 788–791.

[195] C. Donolato, "Modeling the effect of dislocations on the minority carrier diffusion length of a semiconductor," *Journal Of Applied Physics*, vol. 84, no. 5, p. 2656, 1998.

[196] M. Rinio, "Untersuchung der prozessabhängigen Ladungsträgerrekombination an Versetzungen in Siliziumsolarzellen," Ph.D. dissertation, Technische Universität Bergakademie Freiberg, Sep. 2004.

[197] F. Secco d'Aragona, "Dislocation etch for (100) planes in silicon," *Journal Of The Electrochemical Society*, vol. 119, no. 7, pp. 948–951, 1972.

[198] M. I. Bertoni, G. Sarau, D. P. Fenning, M. Rinio, V. Rose, J. Maser, and T. Buonassisi, "Nanoprobe-XRF and micro-Raman Studies of Metal Impurity Decoration around Dislocations," in *Proceedings of the 38th IEEE PVSC*, Austin, TX, Jun. 2013.

[199] H. J. Choi, M. I. Bertoni, J. Hofstetter, D. P. Fenning, D. M. Powell, S. Castellanos, and T. Buonassisi, "Dislocation Density Reduction During Impurity Gettering in Multicrystalline Silicon," *IEEE Journal of Photovoltaics*, vol. 3, no. 1, p. 189, 2013.

[200] T. Lafford, J. Villanova, J. Baruchel, B. Fayard, T. N. Tran Thi, S. Dubois, and D. Camel, "Non-destructive Characterization of Defects in c-Si Solar Cells using Synchrotron Radiation," in *Proceedings of the 6th International Workshop on Crystalline Silicon Solar Cells*, Nov. 2012, pp. 1–4.

[201] T. Ervik, M. Kivambe, G. Stokkan, B. Ryningen, and O. Lohne, "High temperature annealing of bent multicrystalline silicon rods," *Acta Materialia*, vol. 60, no. 19, pp. 6762–6769, 2012.

[202] Y. Wang, D. Srolovitz, J. Rickman, and R. Lesar, "Dislocation motion in the presence of diffusing solutes: a computer simulation study," *Acta Materialia*, vol. 48, no. 9, pp. 2163–2175, 2000.

[203] J. Haunschild, M. Glatthaar, M. Demant, J. Nievendick, M. Motzko, S. Rein, and E. R. Weber, "Quality control of as-cut multicrystalline silicon wafers using photoluminescence imaging for solar cell production," *Solar Energy Materials and Solar Cells*, vol. 94, no. 12, pp. 2007–2012, 2010.

[204] T. Buonassisi, M. D. Pickett, A. A. Istratov, E. Sauer, T. C. Lommasson, E. Marstein, T. Pernau, R. F. Clark, S. Narayanan, S. M. Heald, and E. R. Weber, "Interactions between metals and different grain boundary types and their impact on multicrystalline silicon device performance," in *Proceedings of the 4th WCPEC*, Waikoloa, Hawaii, May 2006, pp. 944–947.

[205] Y.-C. Fan, J. Tan, S. P. Phang, and D. Macdonald, "Iron imaging in multicrystalline silicon wafers via photoluminescence," in *Proceedings of the 35th IEEE PVSC*, Honolulu, HI, 2010, pp. 439–442.