

Precipitated iron: A limit on gettering efficacy in multicrystalline silicon

D. P. Fenning, J. Hofstetter, M. I. Bertoni, G. Coletti, B. Lai et al.

Citation: *J. Appl. Phys.* **113**, 044521 (2013); doi: 10.1063/1.4788800

View online: <http://dx.doi.org/10.1063/1.4788800>

View Table of Contents: <http://jap.aip.org/resource/1/JAPIAU/v113/i4>

Published by the [American Institute of Physics](#).

Related Articles

Towards an optimized all lattice-matched InAlAs/InGaAsP/InGaAs multijunction solar cell with efficiency >50%
Appl. Phys. Lett. **102**, 033901 (2013)

Thin-film-based CdTe photovoltaic module characterization: Measurements and energy prediction improvement
Rev. Sci. Instrum. **84**, 015114 (2013)

Enhanced external quantum efficiency in rectangular single nanowire solar cells
Appl. Phys. Lett. **102**, 021101 (2013)

Modelling of GaAsP/InGaAs/GaAs strain-balanced multiple-quantum well solar cells
J. Appl. Phys. **113**, 024512 (2013)

Ethanol-enriched low-pressure chemical vapor deposition ZnO bilayers: Properties and growth—A potential electrode for thin film solar cells
J. Appl. Phys. **113**, 024908 (2013)

Additional information on *J. Appl. Phys.*

Journal Homepage: <http://jap.aip.org/>

Journal Information: http://jap.aip.org/about/about_the_journal

Top downloads: http://jap.aip.org/features/most_downloaded

Information for Authors: <http://jap.aip.org/authors>

ADVERTISEMENT



AIP Advances

Now Indexed in
Thomson Reuters
Databases

Explore AIP's open access journal:

- Rapid publication
- Article-level metrics
- Post-publication rating and commenting

Precipitated iron: A limit on gettering efficacy in multicrystalline silicon

D. P. Fenning,^{1,a)} J. Hofstetter,¹ M. I. Bertoni,^{1,b)} G. Coletti,² B. Lai,³ C. del Cañizo,⁴ and T. Buonassisi¹

¹*Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

²*ECN Solar Energy, Westerduinweg 3, NL-1755 LE Petten, The Netherlands*

³*Advanced Photon Source, Argonne National Laboratory, Argonne, Illinois 60439, USA*

⁴*Instituto de Energía Solar, Universidad Politécnica de Madrid, 28040 Madrid, Spain*

(Received 6 November 2012; accepted 4 January 2013; published online 31 January 2013)

A 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. © 2013 American Institute of Physics. [<http://dx.doi.org/10.1063/1.4788800>]

I. INTRODUCTION

Semiconductor materials and devices are highly sensitive to dilute concentrations of metal impurities.^{1,2} Silicon-based photovoltaic devices are no exception, with conversion efficiencies typically decreasing as metal impurities exceed atomic concentrations of parts per billion.^{3,4} Iron, in particular, limits the bulk minority carrier lifetime of most as-grown *p*-type silicon wafers,^{5,6} in part because of its large electron capture cross section,^{7,8} but also because of its inevitable presence in feedstocks,⁹ crystal growth crucibles and their linings,¹⁰ and throughout the industrial growth environment.¹¹ Recently, several authors have investigated the macroscopic device-level effects of iron contamination,^{12–14} updating foundational studies of metal contamination in silicon solar cells (e.g., the Westinghouse study of Davis, Jr. *et al.*³) to assess the impact of advances in substrate material quality, higher-efficiency cell architectures, and improved processing.

However, device performance outcomes remain strongly coupled to the execution of defect-engineering techniques during cell processing.^{15,16} 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,^{9,17} by phosphorus diffusion,^{18,19} by extended anneals at temperatures where defects remain

mobile (i.e., low-temperature annealing),^{20–25} and even during metallization firing.^{26,27}

In this contribution, we present a cohesive assessment 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) model²⁸ 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.

II. 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

^{a)}Electronic mail: dfenning@alum.mit.edu.

^{b)}Present address: Arizona State University, Tempe, Arizona 85281, USA.

interstitial iron, Fe_i , and on the concentration of precipitated iron, Fe_p , in form of β -iron-silicide, β -FeSi₂. 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.²⁹

The gettering efficiency for both Fe_i and Fe_p reduction is limited by the solid solubility and the diffusivity of Fe_i in Si, both of which have an Arrhenius relationship with temperature. Consequently, the dissolution of β -FeSi₂ precipitates increases with increasing temperature. However, the segregation coefficient in the gettering layer, which determines the driving force for dissolved Fe_i atoms to diffuse and segregate to the gettering layer, decreases with increasing temperature. The combination of a stronger dissolution of precipitates and lower segregation at higher temperature can result in an increased post-processed Fe_i concentration depending on how the wafers are cooled from high temperature.^{30,31} With the help of the I2E simulation tool,^{28,32} we are able to evaluate these tradeoffs by calculating the post-processed Fe_p and Fe_i concentration for nearly any given time-temperature profile and as-grown iron concentration and distribution.

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 applied to process the wafers used in this study, which originate from the European Integrated Project CrystalClear as described in Sec. III. 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 -FeSi₂ precipitate radii from 8 to 18 nm.³³ 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.

A. 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.³⁴ In many silicon solar cell materials, the pre-gettered and post-gettered lifetimes are limited by iron point defects.⁵ During high-temperature annealing, interstitial iron concentrations increase as iron-silicide precipitates dissolve. In the presence of a gettering layer, the interplay between Fe_i formation (kinetically limited precipitate dissolution) and removal from the bulk (gettering) determines the final post-processed Fe_i concentration, which ultimately impacts solar cell performance.

In addition to describing iron and phosphorus diffusion, coupled by a semi-empirical iron segregation equation,³⁵ our simulation²⁸ assumes Ham's law³⁶ governs the iron precipitation and dissolution kinetics

$$\frac{dC_i}{dt} = 4\pi N_p r D (C_{eq} - C_i), \quad (1)$$

where C_i is the concentration of interstitial iron, D is the temperature-dependent diffusivity, and C_{eq} is the equilibrium iron solubility. The effect of the precipitated iron distribution on dissolution comes from the dependence in Eq. (1) 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 β -FeSi₂, $V_{Fe/FeSi_2}$, and the volume of a single spherical precipitate of average as-grown radius, r_o , we have more explicitly

$$\frac{dC_i}{dt} = 4\pi \frac{C_p \cdot V_{Fe/FeSi_2}}{\frac{4}{3}\pi r_o^3} \cdot r D_{Fe} (C_{eq} - C_i). \quad (2)$$

Fig. 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. 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 are observed in Fig. 1(a) and are described in short below. Further discussion of these regions can be found elsewhere.^{6,37}

1. Region 1

Starting at as-grown total Fe concentrations $\lesssim 2 \times 10^{14}$ cm⁻³, the post-processed Fe_i concentration decreases moving left in Fig. 1(a) toward smaller as-grown total concentrations. As can be seen from Eq. (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 $\lesssim 2 \times 10^{14}$ 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 Eq. (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 a minimum Fe_i concentration is found for total as-grown iron concentrations of $(2-3) \times 10^{13}$ 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.

2. 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




FIG. 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 dotted lines demarcate three different regions of Fe_i response to a standard phosphorus diffusion gettering process.

post-processed Fe_i concentration $1 \times 10^{12} \text{ cm}^{-3}$ is obtained after standard PDG.

3. Region 3

For an as-grown Fe concentration $\geq 4 \times 10^{14} \text{ cm}^{-3}$, 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,^{9,38} such that intragranular precipitates sites may be much further separated than we assume here, leading to decreased internal gettering interactions.

B. Factors determining precipitated iron concentration after gettering

The vast majority of iron atoms in as-grown material are found in FeSi₂ precipitates. However, because of their low spatial density, FeSi₂ precipitates usually have a smaller direct impact on lifetime than Fe_i. Still, FeSi₂ 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. 1(b). It varies between about 1×10^{11} and $7 \times 10^{14} \text{ cm}^{-3}$ 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 $(2-3) \times 10^{13} \text{ cm}^{-3}$ 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 $1.2 \times 10^{13} \text{ cm}^{-3}$,³⁹ 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.^{29,40} The Fe_i diffusivity in Si at 850 °C is about $9.8 \times 10^{-7} \text{ cm}^2/\text{s}$,⁴¹ 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.

III. MATERIALS AND METHODS

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 have already been published.^{12,42,43} 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 be subsequently referred to as “Bottom,” “Middle,” and “Top.” The samples from the bottom were selected ≥ 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,^{42,43} the iron concentrations for the samples are estimated to be: 3.3×10^{13} , 3×10^{13} , 6×10^{14} for the 53 ppmw ingot and




FIG. 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.

TABLE I. 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}	$\Sigma 3$
	88	6×10^{14}	RA (38.7°)
200 ppmw Fe	30	5×10^{13}	$\Sigma 27b$
	68	1.2×10^{14}	RA (27.2°)
	85	4.4×10^{14}	$\Sigma 3$

5×10^{13} , 1.2×10^{14} , 4.4×10^{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.¹² 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 applied 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. 2.

For each set of sister wafers from the six different locations, a single grain boundary (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 I. The regions of interest were selected from the same $2.5 \text{ cm} \times 2.5 \text{ cm}$ 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.⁴⁴ The potential impact of different GB type on the precipitation behavior is discussed in detail in Appendix B.

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 Argonne National Laboratory was used to collect the μ -XRF measurements,⁴⁵ utilizing its 200 nm full-width half-maximum spot size and taking 220 nm steps during μ -XRF mapping with a 1 s 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 Appendix A.

IV. EXPERIMENTAL RESULTS

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. 3. Fig. 4 shows the corresponding maps for the 200 ppmw ingot. The μ -XRF maps were taken at approximately the same position (within $100 \mu\text{m}$) along the same grain boundary in the sister wafers from a given height. The small changes in position along the




FIG. 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 $220 \text{ nm} \times 220 \text{ nm}$. Iron concentrations are shown in units of $\mu\text{g}/\text{cm}^2$. 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.




FIG. 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.

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.

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 Appendix A. A quantification of precipitate size distribution is presented in Fig. 5 for the 53 ppmw ingot and Fig. 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 indicates the automated noise cutoff value for each sample, details about which can be found in Appendix A. Table II 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.*⁴² 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

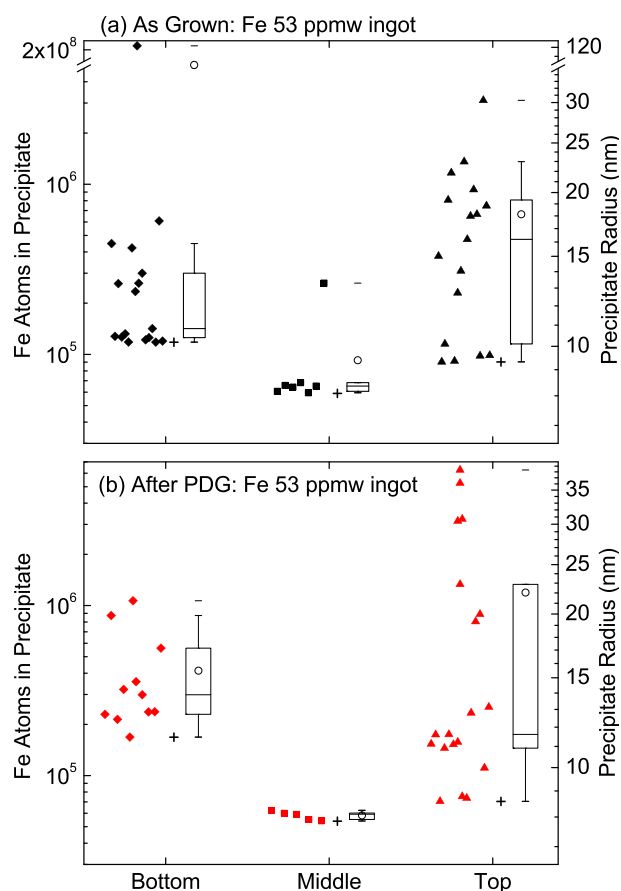


FIG. 5. Distribution of iron in the 53 ppmw ingot (a) before and (b) after gettering. Boxplots show the quartiles of the distribution, and the open circles the mean. The cross underneath the data shows the noise cutoff value from the automated analysis. The bottom and top of the ingot show significant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

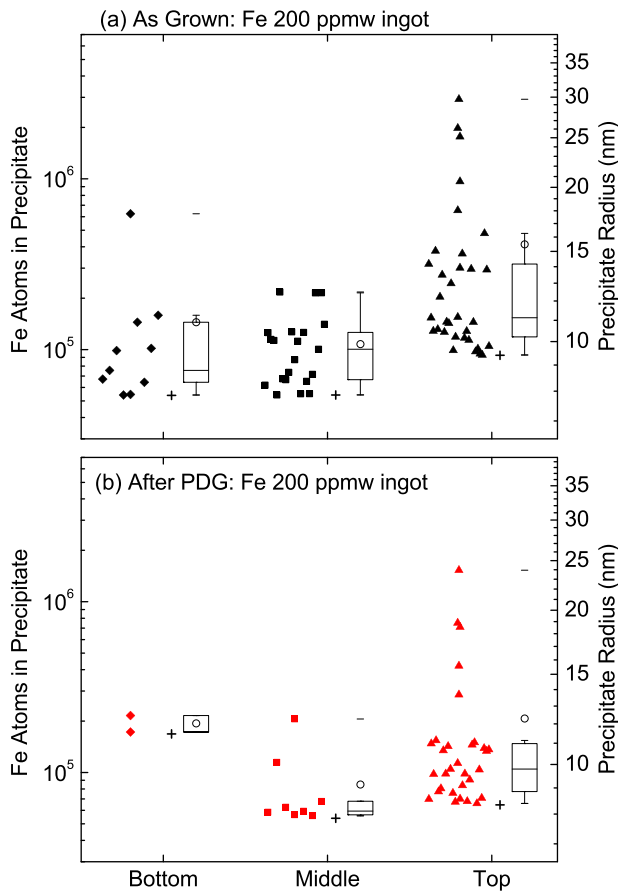


FIG. 6. Distribution of iron in the 200 ppmw ingot (a) before and (b) after gettering. The median as-grown precipitate size increases along the ingot height. The observed precipitate density also increases with ingot height. After phosphorus diffusion, the distributions are shifted toward smaller values at the middle and top of the ingot.

precipitates are near the noise cutoff value, are also close to the fundamental detection limit of our setup, and are not readily observable in Fig. 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

TABLE II. 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

Figs. 5(a) and 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 7×10^4 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. 6. After diffusion, the range of precipitate sizes shifts to smaller values particularly at the top of the ingot as can be seen comparing Figs. 6(a) and 6(b). Additionally, in the phosphorus-diffused middle sample, gettering appears to decrease the number of precipitates larger than $\geq 7 \times 10^4$. After gettering, the bottom sample revealed only a small number of precipitates, in part due to the higher minimum detection limit in that sample as discussed in Appendix A, though the mean size was 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 II. 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.^{26,27} 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. 7. As

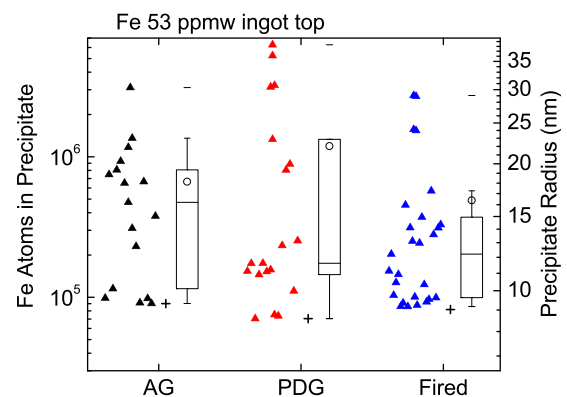


FIG. 7. Direct comparison of the three sister wafers with different processing from the top of the 53 ppmw ingot: as-grown (AG), phosphorus-diffused (PDG), and fired. Significant amounts of iron remain precipitated at grain boundaries after P-diffusion, potentially acting as sources for point defects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

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.

V. DISCUSSION

A. As-grown iron-silicide precipitate size varies 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. 5(a) and 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. 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. 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.*¹⁷ 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.⁴² 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/cm³ and 1×10^{15} atoms/cm³. 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,⁴² 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,⁴⁶ 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.⁹

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.

B. 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 II. A similar trend was found by nucleation and growth simulation by Haarahiltunen *et al.*¹⁷ Schön *et al.*⁴⁷ have also simulated the nucleation and growth of iron-silicide precipitates in wafers with an iron concentration of 3.5×10^{13} atoms/cm³. In such wafers, they expect a much higher density of small precipitates ($10^3 - 5 \times 10^5$ atoms) than large precipitates ($>5 \times 10^4$ atoms). With an experimental detection limit of $\geq 5 \times 10^4$ 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, are 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 Appendix B. 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.

C. 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. 7 summarizes this experimental finding. This lack of significant reduction of precipitated iron is to be expected from the simulations of Fig. 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.^{33,46} 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 reduction 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. 1(b). The as-grown iron distribution measurements in the middle of the 53 ppmw (8 nm median radius, 3×10^{13} atoms/cm³ 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/cm³ total Fe) position it within region 2. Because of the small median radius, the simulations of Fig. 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. 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.*¹² 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 adaptation 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. Similar results indicating little change in iron-silicide precipitates during extended low-temperature annealing have been found previously.^{33,48}

D. 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.^{12,42,43}

Using the same gettering time-temperature profile as in Sec. II, the simulated post-processed interstitial iron concentration is shown for the 53 ppmw and 200 ppmw ingots in Figs. 8(a) and 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^{12,49} 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.




FIG. 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),^{42,43} the as-grown Fe_i concentration (full triangles), and the post-gettering Fe_i concentration (open diamonds).^{12,43,49} 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.

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. 1(a)), 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. 8) and concluded that the high interstitial iron concentrations resulted from poor gettering in these top regions, ultimately leading to the degraded solar cell performance seen there.^{12,43} 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. 1). In the end, solar cells produced from the middle of the 200 ppmw ingot were also iron-limited.⁴³

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. 8. The experimental interstitial iron concentrations after gettering in the middle of the ingot from Ref. 12 approach 10^{11} atoms/cm³, 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 FeSi₂ precipitate was found in the as-grown sample (see Fig. 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 Sec. II, the Fe_i data of Coletti *et al.*¹² shown here plateau 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.*⁴² 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 Sec. II. For high total iron concentrations, the dissolution of FeSi₂ precipitates during standard phosphorus diffusion gettering strongly influences final interstitial iron concentration.

E. 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/cm³. Carrier recombination at the precipitates themselves might limit minority carrier lifetime after PDG as indicated by simulations^{6,47} and recent experimental results.⁴⁷ 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 precipitated 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.^{16,26,27} 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.²⁶

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.*⁵⁰ for aluminum gettering and recently confirmed experimentally for PDG,^{47,51,52} 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:⁵³ for Si ribbon materials, containing a higher density of smaller precipitates due to fast cooling after crystallization,¹¹ 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^{54,55} with customized solar cell processing may allow manufacturers to narrow the standard deviation of solar cell efficiencies while increasing average values.

VI. 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/cm³. 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 the 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.

ACKNOWLEDGMENTS

This work was supported by the U. S. Department of Energy, Contract No. DE-FG36-09GO19001, the MIT-Spain/La Cambra de Barcelona Seed Fund, and the Spanish *Ministerio de Ciencia e Innovación* through Thincells Project No. TEC2008-06798-C03-02 and Crysthin Project No. TEC2011-28423-C03-02. This material is based upon work supported

in part by the National Science Foundation (NSF) and the Department of Energy (DOE) under NSF CA No. EEC-1041895. D. P. Fenning acknowledges the support of the NSF Graduate Research Fellowship. J. Hofstetter acknowledges the support of the Alexander von Humboldt foundation through a Feodor Lynen postdoctoral fellowship. Use of the Advanced Photon Source at Argonne National Laboratory was supported by the U. S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357. EBSD was performed at the Harvard University Center for Nanoscale Systems (CNS), a member of the National Nanotechnology Infrastructure Network (NNIN), which is supported by the National Science Foundation under NSF Award No. ECS-0335765.

APPENDIX A: DETAILS OF μ -XRF DATA ANALYSIS

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 $K\alpha$ absorption edge. The phosphorus-diffused sister wafer at the bottom of both ingots was measured at a subsequent run, and due to experimental constraints a 10.0 keV incident X-ray beam was used. The corresponding effective $1/e$ attenuation length of the fluorescent Fe signal in the μ -XRF setup was about $7 \mu\text{m}$ 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 to convert fluorescence counts into iron concentrations, using a peak fitting procedure by Vogt *et al.*^{56,57} 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.⁵⁸

Following peak-fitting, analysis of the precipitate size distribution requires distinguishing the high-count Fe $K\alpha$ 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. A cutoff concentration was then taken to be the larger of the minimum detection limit or $\mu + 3.5\sigma$, taking pixels with Fe counts below this as noise—pixels with iron concentrations above this cutoff were identified as containing precipitates. 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.

An 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 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 in Figs. 5–7 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.

APPENDIX B: ASSESSMENT OF ASSUMPTIONS FOR COMPARISONS OF EXPERIMENTAL DATA

Comparisons between the as-grown and processed sister wafers must be considered with some caution, because not only is the process condition changing, but the as-grown iron distribution can change as well. 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.⁵⁹ 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 the present study, 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,⁵⁹ which could help explain the smaller variation here.

The GB character measured at the different heights also varied because of the significant changes in ingot height, potentially affecting the measured precipitate distributions. In a study of the dependence of metal contamination on GB character, Buonassisi *et al.*⁴⁴ 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 Sekiguchi⁶⁰ 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.*,⁴⁴ they found that higher order CSL GB character and higher metal contamination lead to higher EBIC contrast. Bertoni *et al.*⁶¹ 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.*⁶² used μ -XRF with 2 μ m 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 Σ 3 GB investigated in that initial study. More recently, Takahashi *et al.*⁶³ used photoluminescence contrast before and after gettering to evaluate the impact GB type has on gettering efficacy. They found that contrast between Σ 3 grain boundaries and the bulk grain generally decreased during gettering, suggesting that Σ 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.

In this study, 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. 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.

¹H. J. Queisser and E. E. Haller, *Science* **281**, 945 (1998).

²A. A. Istratov, H. Hieslmair, and E. R. Weber, *Appl. Phys. A: Mater. Sci. Process.* **70**, 489 (2000).

³J. R. Davis, Jr., A. Rohatgi, R. H. Hopkins, P. D. Blais, P. Rai-Choudhury, J. R. McCormick, and H. C. Mollenkopf, *IEEE Trans. Electron Devices* **27**, 677 (1980).

⁴G. Coletti, P. C. P. Bronsveld, G. Hahn, W. Warta, D. Macdonald, B. Ceccaroli, K. Wambach, N. Le Quang, and J. M. Fernandez, *Adv. Funct. Mater.* **21**, 879 (2011).

⁵B. Sopori, *J. Electron. Mater.* **31**, 972 (2002).

⁶J. Hofstetter, J.-F. Lelièvre, D. P. Fenning, M. I. Bertoni, T. Buonassisi, and C. del Cañizo, *Solid State Phenom.* **178–179**, 158 (2011).

⁷D. Macdonald, A. Cuevas, and J. Wong-Leung, *J. Appl. Phys.* **89**, 7932 (2001).

⁸S. Rein and S. W. Glunz, *J. Appl. Phys.* **98**, 113711 (2005).

⁹T. 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. Saunar, and E. R. Weber, *Prog. Photovoltaics* **14**, 513 (2006).

¹⁰E. Olsen and E. J. Øvrelid, *Prog. Photovoltaics* **16**, 93 (2008).

¹¹T. Buonassisi, A. A. Istratov, M. Heuer, M. A. Marcus, R. Jonczyk, J. Isenberger, B. Lai, Z. Cai, S. Heald, W. Warta, R. Schindler, G. Willeke, and E. R. Weber, *J. Appl. Phys.* **97**, 074901 (2005).

¹²G. Coletti, R. Kvande, V. D. Mihailetchi, L. J. Geerligs, L. Arnberg, and E. J. Øvrelid, *J. Appl. Phys.* **104**, 104913 (2008).

¹³S. Riepe, I. E. Reis, W. Kwapil, M. A. Falkenberg, J. Schön, H. Behnken, J. Bauer, D. Kreßner-Kiel, W. Seifert, and W. Koch, *Phys. Status Solidi C* **8**, 733 (2011).

¹⁴S. Dubois, O. Palais, P. J. Ribeyron, N. Enjalbert, M. Pasquinelli, and S. Martinuzzi, *J. Appl. Phys.* **102**, 083525 (2007).

- ¹⁵S. A. McHugo, H. Hieslmair, and E. R. Weber, *Appl. Phys. A: Mater. Sci. Process.* **64**, 127 (1997).
- ¹⁶J. Tan, D. Macdonald, N. Bennett, D. Kong, A. Cuevas, and I. Romijn, *Appl. Phys. Lett.* **91**, 043505 (2007).
- ¹⁷A. Haarahiltunen, H. Savin, M. Yli-Koski, H. Talvitie, M. I. Asghar, and J. Sinkkonen, *Mater. Sci. Eng., B* **159**, 248 (2009).
- ¹⁸D. Macdonald, A. Cuevas, A. Kinomura, and Y. Nakano, in *Proceedings of the 29th IEEE PVSC* (New Orleans, LA, 2002), pp. 285–288.
- ¹⁹M. Seibt, A. Sattler, C. Rudolf, O. Voß, V. Kveder, and W. Schröter, *Phys. Status Solidi A* **203**, 696 (2006).
- ²⁰M. D. Pickett and T. Buonassisi, *Appl. Phys. Lett.* **92**, 122103 (2008).
- ²¹J. Härkönen, V.-P. Lempinen, T. Juvonen, and J. Kylmäluoma, *Sol. Energy Mater. Sol. Cells* **73**, 125 (2002).
- ²²T. Buonassisi, A. A. Istratov, M. A. Marcus, B. Lai, Z. Cai, S. M. Heald, and E. R. Weber, *Nature Mater.* **4**, 676 (2005).
- ²³P. Manshanden and L. J. Geerligts, *Sol. Energy Mater. Sol. Cells* **90**, 998 (2006).
- ²⁴R. Krain, S. Herlufsen, and J. Schmidt, *Appl. Phys. Lett.* **93**, 152108 (2008).
- ²⁵M. Rinio, A. Yodyunyong, S. Keipert-Colberg, Y. P. B. Mouafi, D. Borchert, and A. Montesdeoca-Santana, *Prog. Photovoltaics* **19**, 165 (2011).
- ²⁶J.-F. Lelièvre, J. Hofstetter, A. Peral, I. Hoces, F. Recart, and C. del Cañizo, *Energy Procedia* **8**, 257 (2011).
- ²⁷B. Michl, J. Schön, F. Schindler, W. Warta, and M. C. Schubert, in *27th Proceedings of the EUPVSEC* (Frankfurt, Germany, 2012), pp. 709–713.
- ²⁸J. Hofstetter, D. P. Fenning, M. I. Bertoni, J.-F. Lelièvre, C. del Cañizo, and T. Buonassisi, *Prog. Photovoltaics* **19**, 487 (2011).
- ²⁹D. Macdonald, A. Cuevas, A. Kinomura, Y. Nakano, and L. J. Geerligts, *J. Appl. Phys.* **97**, 033523 (2005).
- ³⁰C. Ballif, S. Peters, D. Borchert, C. Hässler, J. Isenberg, R. Schindler, W. Warta, and G. Willeke, in *Proceedings of the 17th EU PVSEC* (Munich, Germany, 2001), p. 1818.
- ³¹S. Peters, J. Y. Lee, C. Ballif, D. Borchert, S. W. Glunz, W. Warta, and G. Willeke, in *Proceedings of the 29th IEEE PVSC* (New Orleans, LA, 2002), pp. 214–217.
- ³²See <http://pv-i2e.mit.edu> for Impurities-to-Efficiency (I2E) Simulator.
- ³³D. P. Fenning, J. Hofstetter, M. I. Bertoni, S. Hudelson, M. Rinio, J. F. Lelièvre, B. Lai, C. del Cañizo, and T. Buonassisi, *Appl. Phys. Lett.* **98**, 162103 (2011).
- ³⁴A. A. Istratov, H. Hieslmair, and E. R. Weber, *Appl. Phys. A: Mater. Sci. Process.* **69**, 13 (1999).
- ³⁵A. Haarahiltunen, H. Savin, M. Yli-Koski, H. Talvitie, and J. Sinkkonen, *J. Appl. Phys.* **105**, 023510 (2009).
- ³⁶F. Ham, *J. Phys. Chem. Solids* **6**, 335 (1958).
- ³⁷J. Hofstetter, “Defect engineering strategies for solar grade silicon and their optimization by predictive simulation,” Ph.D. dissertation (ETSI Telecomunicación (Universidad Politécnica de Madrid, 2011).
- ³⁸M. I. Bertoni, D. P. Fenning, M. Rinio, V. Rose, M. Holt, J. Maser, and T. Buonassisi, *Energy Environ. Sci.* **4**, 4252 (2011).
- ³⁹M. Aoki, A. Hara, and A. Ohsawa, *J. Appl. Phys.* **72**, 895 (1992).
- ⁴⁰T. U. Naerland, L. Amberg, and A. Holt, *Prog. Photovoltaics* **17**, 289 (2009).
- ⁴¹E. R. Weber, *Appl. Phys. A: Mater. Sci. Process.* **30**, 1 (1983).
- ⁴²R. Kvande, L. J. Geerligts, G. Coletti, L. Arnberg, M. Di Sabatino, E. J. Øvrelid, and C. C. Swanson, *J. Appl. Phys.* **104**, 064905 (2008).
- ⁴³G. Coletti, P. C. P. Bronsveld, R. Kvande, L. J. Geerligts, and C. C. Swanson, in *Arriving at Well-Founded Specifications of Solar Grade Silicon* (Amsterdam, The Netherlands, 2008).
- ⁴⁴T. Buonassisi, A. A. Istratov, M. D. Pickett, M. A. Marcus, T. F. Ciszek, and E. R. Weber, *Appl. Phys. Lett.* **89**, 042102 (2006).
- ⁴⁵Z. 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, in *AIP CP 521, Synchrotron Radiation Instrumentation: Eleventh US National Conference* (Stanford, CA, 1999), pp. 31–34.
- ⁴⁶J. Schön, A. Haarahiltunen, H. Savin, D. P. Fenning, T. Buonassisi, W. Warta, and M. C. Schubert, *IEEE J. Photovoltaics* **3**, 131 (2011).
- ⁴⁷J. Schön, H. Habenicht, M. C. Schubert, and W. Warta, *J. Appl. Phys.* **109**, 063717 (2011).
- ⁴⁸M. Rinio, A. Yodyunyong, S. Keipert-Colberg, D. Borchert, and A. Montesdeoca-Santana, *Phys. Status Solidi A* **208**, 760 (2011).
- ⁴⁹G. Coletti, private communication (2012).
- ⁵⁰P. S. Plekhanov, R. Gafiteanu, U. M. Gösele, and T. Y. Tan, *J. Appl. Phys.* **86**, 2453 (1999).
- ⁵¹V. Osinniy, A. Nylandsted Larsen, E. Hvidsten Dahl, E. Enebakk, A.-K. Søiland, R. Tronstad, and Y. Safir, *Sol. Energy Mater. Sol. Cells* **101**, 123 (2012).
- ⁵²D. P. Fenning, A. S. Zuschlag, M. I. Bertoni, B. Lai, G. Hahn, and T. Buonassisi, “Improved Iron Gettering of Contaminated Multicrystalline Silicon by High-Temperature Phosphorus Diffusion”, *Progress in Photovoltaics: Research and Application* (submitted).
- ⁵³J. Hofstetter, D. P. Fenning, J.-F. Lelièvre, C. del Cañizo, and T. Buonassisi, *Phys. Status Solidi A* **209**, 1861 (2012).
- ⁵⁴T. Mankad, R. Sinton, A. Ristow, N. Akil, and E. Picard, in *21st Workshop on Crystalline Silicon Solar Cells and Modules* (Breckenridge, CO, 2011).
- ⁵⁵T. Trupke, K. R. McIntosh, J. W. Weber, W. McMillan, L. Ryves, J. Haunschild, C. Shen, and H. Kampwerth, in *22nd Workshop on Crystalline Silicon Solar Cells and Modules* (Vail, CO, 2012).
- ⁵⁶S. Vogt, J. Maser, and C. Jacobsen, *J. Phys. IV* **104**, 617 (2003).
- ⁵⁷S. Vogt, *J. Phys. IV* **104**, 635 (2003).
- ⁵⁸A. Zuschlag and G. Hahn, in *Proceedings of the 26th EU PVSEC* (Hamburg, Germany, 2011), pp. 1547–1549.
- ⁵⁹A. Zuschlag, S. Ohl, J. Bernhard, H. Morhenn, J. Ebser, J. Junge, S. Seren, and G. Hahn, in *Proceedings of the 35th IEEE PVSC* (Honolulu, HI, 2010), pp. 347–351.
- ⁶⁰J. Chen and T. Sekiguchi, *Jpn. J. Appl. Phys., Part 1* **46**, 6489 (2007).
- ⁶¹M. 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, *Prog. Photovoltaics* **19**, 187 (2011).
- ⁶²K. Arafune, E. Ohishi, H. Sai, Y. Terada, Y. Ohshita, and M. Yamaguchi, *Jpn. J. Appl. Phys., Part 1* **45**, 6153 (2006).
- ⁶³I. Takahashi, N. Usami, H. Mizuseki, Y. Kawazoe, G. Stokkan, and K. Nakajima, *J. Appl. Phys.* **109**, 033504 (2011).