Solar Energy Materials & Solar Cells 114 (2013) 110–135

Contents lists available at SciVerse ScienceDirect

Solar Energy Materials & Solar Cells

journal homepage: www.elsevier.com/locate/solmat

A wafer-based monocrystalline silicon photovoltaics road map: Utilizing known technology improvement opportunities for further reductions in manufacturing costs

Alan Goodrich*, Peter Hacke, Qi Wang, Bhushan Sopori, Robert Margolis, Ted L. James, Michael Woodhouse**

The National Renewable Energy Laboratory, Golden, CO USA

ARTICLE INFO

Article history:
Received 20 July 2012
Received in revised form
10 January 2013
Accepted 22 January 2013

Keywords:
Crystalline silicon
Photovoltaics
Solar energy
Economics

ABSTRACT

As an initial investigation into the current and potential economics of one of today's most widely deployed photovoltaic technologies, we have engaged in a detailed analysis of manufacturing costs for each step within the wafer-based monocrystalline silicon (c-Si) PV module supply chain. At each step we find several pathways that could lead to further reductions in manufacturing costs. After aggregating the performance and cost considerations for a series of known technical improvement opportunities, we project a pathway for commercial-production c-Si modules to have typical sunlight power conversion efficiencies of 19–23%, and we calculate that they might be sustainably sold at ex-factory gate prices of \$0.60–\$0.70 per peak Watt (DC power, current U.S. dollars).

This may not be the lower bound to the cost curve for c-Si, however, because the roadmap described in this paper is constrained by the boundary conditions set by the wire sawing of wafers and their incorporation into manufacturing equipment that is currently being developed for commercial-scale production. Within these boundary conditions, we find that the benefit of reducing the wafer thickness from today's standard 180 μm to the handling limit of 80 μm could be around \$0.05 per peak Watt (W_p), when the calculation is run at minimum sustainable polysilicon prices (which we calculate to be around \$23/kg). At that minimum sustainable polysilicon price, we also calculate that the benefit of completely eliminating or completely recycling kerf loss could be up to \$0.08/W_p.

These downward adjustments to the long run wafer price are used within the cost projections for three advanced cell architectures beyond today's standard c-Si solar cell. Presumably, the higher efficiency cells that are profiled must be built upon a foundation of higher quality starting wafers. The prevailing conventional wisdom is that this should add cost at the ingot and wafering step—either due to lower production yields when having to sell wafers that are doped with an alternative element other than the standard choice of boron, or in additional capital equipment costs associated with removing problematic boron–oxygen pairs. However, from our survey it appears that there does not necessarily need to be an assumption of a higher wafer price if cell manufacturers should wish to use *n*-type wafers derived from the phosphorus dopant. And as for making *p*-type wafers with the traditional boron dopant, the potential price premium for higher lifetimes via the magnetic *Czochralski* approach is calculated to be very small, and can ostensibly be offset by the higher expected cell efficiencies that would result from using the higher quality wafers. With this final consideration, the projected minimum sustainable price requirements for three advanced c-Si solar cells are incorporated into a final bill of materials for a polysilicon-to-module manufacturing facility located within the United States.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

With average annual growth rates in excess of 40% over the past decade [1,2], the success of the PV industry can largely be attributed to the steadfast growth of wafer-based multicrystalline and monocrystalline silicon. This growth has been sustained through a powerful combination of three critical competitive advantages: (1) industry-leading full module area sunlight power conversion efficiencies (to date, monocrystalline silicon continues to provide the highest power conversion efficiency among all commercially demonstrated single junction PV modules [3,4]); (2) product 'bankability' from the appropriately qualified suppliers (with warranties for 80% of original performance after 25 years

* Corresponding author. Tel.: +1 303 275 4347.
** Corresponding author. Tel.: +1 720 883 4973.
E-mail addresses: Alan.Goodrich@nrel.gov (A. Goodrich), Michael.Woodhouse@nrel.gov (M. Woodhouse).

0927-0248/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.solmat.2013.01.030

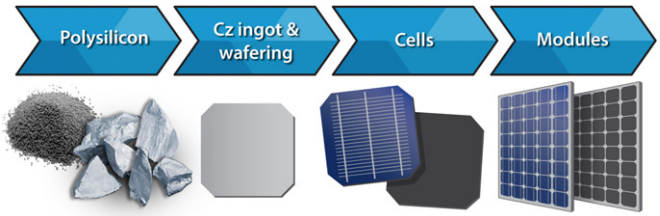

Fig. 1. The primary steps of the wafer-based c-Si module supply chain.

of service now being standard [5]); and (3) a consistent ability to offer competitively priced modules, which has been enabled through an ability to realize cost reductions throughout the c-Si module supply chain (see Fig. 1).

A significant portion of these cost reductions have come about due to 'economies-of-scale' benefits [6]. But there is a point of diminishing returns when trying to lower costs by simply expanding production capacity. For PV to sustain its trajectory of cost reductions in both manufacturing and systems, it will be increasingly important to implement innovations that enable higher sunlight power conversion efficiencies [7]; and while the advanced cell architectures needed to achieve these higher efficiencies require a greater initial capital expenditure in the manufacturing equipment and starting materials, sufficient gains in efficiency can oftentimes work to offset these added costs. For wafer-based c-Si there are also multiple pathways to lower costs further through reductions in the cost of producing the polysilicon feedstock, better silicon utilization in wafer fabrication, and through advances in industrial cell and module assembly processes.

Separate from specific technology advancements, there are also pathways to lower future costs if an industry-wide supply-demand equilibrium can be reached. The prices for all materials within the supply chain could even approach their minimum sustainable levels, at the point of a perfectly balanced equilibrium. Finally, a 'vertical integration' strategy—where the buyers and sellers in the supply chain are united into a single firm or consortium—can also assist in lowering material transfer prices. Driven by differences in technology focus, regional differences in electricity and labor rates, and its still small scale relative to more mature industries, the supply chain for c-Si modules has historically been comprised of distinct firms specializing in polysilicon feedstock, wafers, cells or modules. Most recently, however, the vertical integration strategy has come to play an increasingly evident role, because it can provide several significant competitive advantages.

The principal advantage of global supply-demand equilibrium and the vertical integration strategy is that they enable better control over the often volatile, market-driven price demands of upstream suppliers. As an unambiguous demonstration of this point, one can consider the recent trends and effects of polysilicon prices. The average spot price that wafer manufacturers had to pay for this material rapidly rose from around $200/kg in 2007 to highs around $400/kg in 2008, principally because poly suppliers could readily command the higher price (due, in no small part, to the sudden increase in demand from so many rapidly emerging PV companies [8]). Since those times, new polysilicon factories have come online at a frenetic pace, and the resulting current oversupply situation has forced polysilicon suppliers to lower spot prices to less than $20/kg within the past year [9]. Meanwhile, global average c-Si module prices have also recently taken a plunge—from around $4/W in 2008 to less than $1/W today [10,11]. Given how important the decrease in the price for poly has been to realizing those dramatic decreases in the total module price, there is certainly a precedent that—for each player in the supply chain—minimum sustainable transfer pricing could be one key strategy to survival in a game that will be won or lost by pennies-per-watt.

For many, it seems compelling that PV can provide such valuable contributions to the dual challenges of carbon mitigation and overall energy security. But, in addition to the integration challenges associated with its intermittent nature, there are also a number of economic barriers that must still be overcome before there can be energy-significant Terawatts adoption levels. Primary among these economic barriers is that, in most cases, PV is still an overall relatively expensive choice for power generation.

With the recent drops in module and system prices, such economic barriers are becoming that much less daunting. Under certain conditions (specifically, in markets where there are high traditional fuel costs and relatively intense solar irradiation levels), the barrier of a higher levelized cost for generating electrical power from PV, in comparison to traditional power generation sources, has even recently been jumped [11–14]. In looking forward, to understand how PV might go on to become a compelling choice under other conditions, it becomes necessary to understand just how it compares to the incumbent energy systems in other applications and locations. With the general paradigm being that PV must continue to reduce costs even further, in order to become more economically competitive, it therefore becomes critical to understand just how low the price for complete systems can be without compromising the financial viability of all the different players that are involved.

Though it is beyond the scope of this paper to address total system costs, and to suggest a final LCOE number, in the pages below we shall endeavor to derive the lower limit in price for the module component of PV system costs—for the specific case of c-Si modules made within the United States from *Czochralski*-grown, wire-sawn wafers. For any complete product, its minimum sustainable price is bound by the sum of the minimum sustainable prices for every upstream material in the product's supply chain. Conceptually speaking, such a lower limit would only be realized once the supply-demand equilibrium dictates that all materials be sold precisely at their minimum sustainable prices. Assuming such a scenario, in this paper we shall show how the cost and price requirements for c-Si modules could evolve over time, after detailing the cost parameters associated with several representative technology 'roadmaps'. We begin with an overview of our modeling methodology, and describe the methods and assumptions that are used for calculating the minimum sustainable prices of each node within the c-Si module supply chain.

2. Methods for establishing the minimum sustainable product prices

For each of the primary steps within the c-Si supply chain, the structure of our Microsoft Excel-based cost models is built around the process flow that would be relevant to the manufacturing process being considered. These process flows are based upon extensive literature surveys, internal discussions with NREL researchers, visits to facilities already in place, and through extensive collaborations with PV researchers and company representatives. With the process flows in hand, we then aggregate the typical manufacturing cost considerations for each underlying step, with data that is provided by several industry collaborators that are involved at each step. These considerations include the relevant materials and manufacturing equipment costs; operational costs (which can be calculated after knowing the labor requirements for each piece of equipment, material yield losses, the total cycle times for each step, utilities costs, etc.); and the typical costs for financing the initial capital expenditure in the land, building, and equipment [15]. The specific operational costs of labor and electricity are

estimated by applying national average rates [16] to the staffing and power requirements provided by the relevant equipment vendors.

After considering the expected throughput, maintenance downtime, etc. for each piece of equipment needed to execute the model process flow, the total capital equipment needs are calculated to meet an annual production volume set to meet the full economies-of-scale benefits for that specific node in the supply chain (which is, by our analysis of costs as a function of production volume, around 15,000 metric tonnes per year for polysilicon production, tens of millions of wafers per month in the pulling of ingots and wafering, and 500 $MW_{p(DC)}$–2 $GW_{p(DC)}$ of annual production for cells and modules). In addition to purchasing the capital equipment necessary to begin production, the other upfront capital expenditure that must be included is the cost for the manufacturing facility, including all of the buildings and land necessary for housing the manufacturing project. We calculate these facilities-related expenses in consultation with the equipment vendors and in consultation with companies already engaged in the related manufacturing activities.

After deriving the total initial capital expenditure for the equipment and facilities, the first derived cost metric of interest is the manufacturing 'CapEx', which is calculated by dividing the total initial capital expenditure by the expected production volume of the model manufacturing facility. In this way our CapEx costs are expressed on a dollars-per-kg (polysilicon), dollars-per-area (wafers), or dollars-per-watt (cells and modules) basis.

In any typical manufacturing project, however, the total initial capital expenditure can be shared over a certain length of time—a 'depreciation period'—that is set by country-specific tax codes and by an assumed rate of technology obsolescence. In consideration of these factors, to represent the averaged depreciation costs we allocate the total initial capital equipment expenditure over a ten-year, straight-line depreciation schedule in polysilicon and wafering; a five-year, straight-line schedule in cells; and a seven-year, straight-line schedule in modules. The average annual depreciation expense for the building is also calculated by allocating the initial capital expenditure over a linear depreciation schedule, with an assumed period of 30 years. By dividing these so-allocated depreciation expenditures by the specified production volumes, we are then able to represent a second cost metric of interest—the average depreciation cost—in the same units as the CapEx.

Within the final section of our Excel-based models we set up a *pro forma* discounted cash flow (DCF) for the model manufacturing facility. The purpose of the DCF is to provide the necessary framework for deriving the minimum sustainable prices for each product within the supply chain. Within the DCF, we are able to account for several additional considerations for manufacturing, such as inflation and taxes; typical sales, general and administrative (SG & A) expenses; typical research and development (R & D) expenses; and warranty coverage [17]. Additionally, because it is a DCF, we have the option of distributing the initial equipment and facilities expenditures over Modified Accelerated Cost Recovery System (MACRS) depreciation schedules (with assumed depreciation periods being the same as those given above for each respective step in the supply chain). For each product in the c-Si supply chain, the minimum sustainable price is then calculated on the basis of attaching a minimum required margin, above the nominal cost of production, to satisfy the returns on investment that would be required from both debt investors (such as banks) and equity investors (such as globally-distributed stockholders). With the total length of the DCF set by the length of the assumed depreciation schedule, and the discount rate calculated from these required rates of return, the minimum sustainable product price is then derived by the iterative algorithm within Excel called 'goal seek', which runs until the net present value of the free cash flows equals the total initial capital expenditure.

For any given manufacturing firm, the required rate of return from investors is usually derived after carrying out an assessment of the firm's risk relative to other investment opportunities. Such an assessment can even be quantified after knowing the firm's equity 'beta' (a measure of correlation between the performance of the firm's stock and the performance of a general global investment index, such as Standard and Poor's 500); and after knowing the firm's capital financing structure, which consists of both the book value of debt (BV_D) and the market value of equity (MV_E). These risk factors then influence the rates of return required for both debt financing (K_D) and for attracting equity financing (K_E). We use the international capital assets pricing model to derive these debt and equity rates, and then, by weighting them by their relative contribution to the overall capital structure of the firm, arrive at a weighted average cost of capital (WACC) [18]

$$\mathrm{WACC} = \left(\frac{MV_E}{BV_D + MV_E}\right) K_E + \left(\frac{BV_D}{BV_D + MV_E}\right) K_D (1-T) \tag{1}$$

Within this expression, the 'leverage ratio' is the relative amount of debt (i.e., the $BV_D/(BV_D+MV_E)$ term), and T is the corporate income tax rate. This tax rate, and the host of risk factors that influence the expected rates for debt and equity financing, is heavily dependent upon which country hosts the manufacturing project. For the sake of brevity, within this paper we limit the estimated WACC (Table 1), leverage ratios, and all other representative PV manufacturing costs to what would be typical for manufacturing within the United States, where the average *effective* corporate tax rate for publicly traded, non-financial institutions is around 28% [19]. An additional discussion of how these rates, leverage ratios, and taxes might vary for an alternative manufacturing site is provided elsewhere [20,21].

If a firm is publicly traded, all of the inputs needed for a WACC calculation are updated daily and are available online [22]. After carrying out the relevant calculations for several noteworthy U.S.-based players in the c-Si supply chain, we estimate that a WACC of around 8.6% would have been representative for the first half of 2012.

For several reasons, however, the inputs for a WACC calculation on PV manufacturing could change over time. For example, if the prices for modules and systems continue to fall so as to mitigate the industry's dependence upon subsidies, and if utilities more widely adopt PV systems because they view them to be a sensible substitute for their usual choices, then one might expect that the risks associated with investments into PV companies and installation projects will eventually become more like the risks associated with investments into traditional sources for power generation. So assuming that the recent return requirements of several conventional energy companies might also be representative of what U.S.

Table 1
WACC assumptions used within this paper for U.S. c-Si PV manufacturing. The First-Half (1H) 2012 WACC is used for the benchmarking technology cases shown throughout this paper; the long-term market WACC assumption will be specified within the figure legends or captions when it is used. As an important note, we use the five-year beta when estimating K_E in the International CAPM methodology.

Weighted average cost of capital (WACC) assumptions used for derivations of minimum sustainable prices

	1H 2012	**Long-term**
$\frac{MV_E}{BV_D+MV_E}$	0.60	0.70
Levered cost of equity (K_E)	12%	7.5%
Leverage ratio $\frac{BV_D}{BV_D+MV_E}$	0.40	0.30
Levered cost of debt (K_D)	4.5%	4.5%
Corporate tax rate (T)	28%	28%
WACC	**8.6**%	**6.2**%

PV might look like over the long-term [23], we also detail the inputs for our long-term WACC case in Table 1.

The WACC estimates derived in Table 1 are used as the discount rates in our *pro forma* approach to estimating minimum sustainable product prices. Because not even the vertical integration strategy can eliminate the need for every material within a product's supply chain to meet a minimum required rate of return [24], our best-case, long-term minimum sustainable price projections for complete c-Si modules are built upon the assumption that the transfer prices for all upstream materials are precisely set to achieve a minimum required margin, on top of the nominal cost of production, which is calculated to meet these WACCs. Thus, the values in Table 1 are crucial assumptions that are compounded within our analysis.

3. Polysilicon feedstock

3.1. The Siemens process for producing polysilicon chunk

The very first step in the fabrication of a c-Si wafer is the production of metallurgical grade silicon via the high-temperature reduction of silica (the source of which is typically lumpy quartz, not sand). With coke serving as the reducing agent, the process is most typically carried out in an electric arc furnace with carbon electrodes [25]

$$SiO_2 + 2C \xrightarrow{1900\ ^\circ C} Si(mg) + 2CO$$

The elemental purity of this metallurgical grade silicon, which currently sells for around $2.50/kg, is approximately 98%. But the material purity requirement for the highest efficiency c-Si devices can approach 99.9999999% (9N). The most widely used process for the production of the much more pure polysilicon feedstock material is a chemical vapor deposition (CVD) method called the *Siemens* process, whose processing sequence is broadly represented in Fig. 2.

In order to remove the impurities contained within metallurgical grade silicon, the first step in the *Siemens* CVD process involves the production and distillation of trichlorosilane (TCS). Facilities that manufacture more than 2000 metric tons per annum (MTPA) of polysilicon generally manufacture their own TCS onsite. The production of TCS can be achieved by the reaction of metallurgical grade silicon with hydrochloric acid at moderate temperatures. Most of the impurities that were present within the metallurgical grade Si are left behind while the TCS is distilled [25]

$$Si(mg) + 3HCl \xrightarrow{300\ ^\circ C} SiHCl_3(l) + H_2$$

Solid polysilicon is then produced in a batch process as TCS is converted over the surface of silicon rods that have been placed

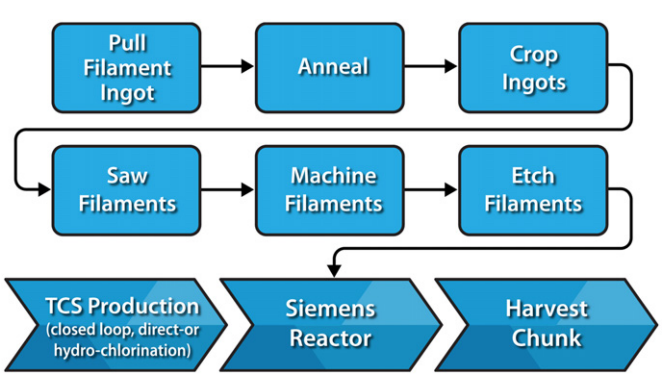

Fig. 2. Generalized process flow for the production of solar grade polysilicon feedstock via the *Siemens* process.

inside of large bell jars, or '*Siemens* reactors' as they are commonly called. These silicon rods—or 'filaments'—are produced from ingots made from either the *Czochralski* (*Cz*) or Float Zone (*FZ*) approaches. The as-produced filaments of today are typically a 7 mm × 7 mm × 2500 mm elongated square, which have been sawn lengthwise from the ingots using slurry-based wire saws. The cropped ingot scrap can be reused for making other ingots, but, due to inclusions of chemical impurities from the wire-sawing slurry, and because it remains in the form of a very fine powder that is extremely difficult to mechanically separate from the SiC based slurry used during the cutting process, the approximately 10–15% of the ingot removed as sawing—or 'kerf'—loss has essentially no value. As final steps before the CVD chamber is sealed, the filaments are mechanically shaped to fit the electrical contacts made for each, a bridge of filament material is set in place between each parallel pair, and the native oxide is etched off using a dilute aqueous HF solution.

Electrical current is passed through the resistive U-shaped silicon filaments to reach a temperature that approaches 1150 °C. This rather high temperature serves to activate the growth of solid polysilicon, Si (*ps*), on the surface of these filaments as a result of the hydrogenation of TCS with an HCl catalyst. The decomposition of trichlorosilane to produce dichlorosilane (SiH_2Cl_2) is one of several side reactions that also occur in the course of this growth process. Fortunately, this intermediate can also react to make polysilicon, and so—even though the TCS stream usually contains 6–9% DCS—most polysilicon producers choose not to separate the two. This leaves the reaction series to be most generally described as follows [25,26]:

$$\begin{array}{lllll} SiHCl_3 & + & H_2 & \xrightarrow{1150\ ^\circ C} & Si\ (ps) + 3HCl \\ 3SiHCl_3 & + & 3HCl & \leftrightarrow & 3SiCl_4 + 3H_2 \\ \hline 4SiHCl_3 & & & \rightarrow & Si\ (ps) + 3SiCl_4 + 2H_2 \end{array}$$

A leading high-pressure 500-MTPA reactor made in 2012 would accommodate 72 rods; the *Siemens* process would typically stopped once a diameter of 125 mm is reached for each. In a reactor of that size, approximately 125 kg of hydrogen is consumed during each hour of polysilicon growth, and the process is approximately 20% efficient in its use of TCS for each pass through the chamber. A total processing time of approximately 60 h per batch is typical, including a total time of around 24 h for filament placement, oxide etching, and for harvesting of the U-shaped polysilicon rods. As final steps, the polysilicon rods are smashed into chunks and packaged in nitrogen- or argon-filled bags for shipping.

In order to drive the reaction sequence toward the production of polysilicon, it is helpful to remove the H_2 and $SiCl_4$ as they are produced within the bell jar. Fortunately, these effluents are actually useful in that they can be recycled for the production of trichlorosilane (which can, of course, be used again in later rounds of polysilicon production). The hydrogenation of silicon tetrachloride, more commonly called the 'direct chlorination' method, is one such pathway

$$H_2 + SiCl_4 \rightarrow SiHCl_3 + HCl$$

Or the H_2 and $SiCl_4$ can be reacted with metallurgical grade Si in the 'hydrochlorination' process [26]

$$2H_2 + 3SiCl_4 + Si(mg) \rightarrow 4SiHCl_3$$

The yields for the hydrochlorination route are generally more difficult to control and it is a more technically challenging process. Thus, those companies having less experience—but also a desire to quickly scale up and establish a presence in this upstream step of the supply chain—are more likely to adopt the

direct chlorination approach [27]. The direct chlorination method does, however, require nearly double the capital equipment investment and uses significantly more energy: 120–200 kWh/kg for direct chlorination versus 65–90 kWh/kg for hydrochlorination [28].

3.2. The fluidized bed reactor (FBR) process for producing polysilicon

The process of polysilicon production via the fluidized bed process is an altogether *physically* different approach from the *Siemens* process [29]. The end product is also quite different in that polysilicon granules, ranging in size from 100 to 1500 μm, are produced instead of the much larger chunks [30].

A fluidized bed reactor is a cone shaped reaction vessel containing small crystalline silicon seed particles that are suspended by an upward-flowing 'fluidizing' gas. This becomes physically possible once the upward drag force of the fluidizing gas is approximately equal to the downward gravitational pull on the particle, based upon its mass ($W = mg$). At the same time they are being fluidized, the particles must be heated above the decomposition temperature of a silicon precursor gas (commonly SiH_4) that is introduced into the vessel. Once the necessary decomposition temperature is reached, with hydrogen serving as the fluidizing gas purified crystalline silicon layers build up layer-upon-layer onto the suspended silicon beads. After reaching a size whereby their weight becomes greater than the upward drag force of the fluidizing gas, the heavier crystallized Si granules fall to the bottom of the cone where they are collected.

There are several advantages to this approach in that it is much more efficient in the overall net use of the reactant gases; it does not require the fabrication, shaping, and placement of crystalline seed filaments; and it requires significantly less energy, at only around 12–20 kWh/kg [8]. The material form factor of the FBR granules is also quite advantageous in the subsequent step of melting polysilicon because the granules can be continuously fed into *Cz* pullers to bear up to 3 daughter ingots per initial charge (versus having to reload polysilicon chunk in single batch processing). In addition, the semi-continuous feeding of granules enables the semi-continuous feeding of dopants; and this can be helpful in overcoming the well-known challenges of uniformly distributing dopants having low segregation coefficients [31]. In spite of its numerous apparent advantages, however, there are also numerous technical challenges in qualifying new FBR facilities. In particular, it can be difficult to manage the heating of the fluidized beads in a controlled manner, without losing an important temperature differential between the reaction zone and the walls of the reactor cone [29]. This at least partially explains why there are currently only a handful of companies that have the technical capability to provide this FBR material.

3.3. Cost model results for polysilicon production

In Fig. 3, the manufacturing cost model results are shown for the two approaches to polysilicon production most commonly employed within the U.S. In both cases, the largest expense is calculated to be the average depreciation on the capital equipment and manufacturing facility (as the basis for our calculations, the total calculated upfront capital expenditure in the equipment and facilities, or 'CapEx', worked out to be $74/kg of annual production capacity for a *Siemens* hydrochlorination facility and $71/kg of annual production capacity for an FBR-based facility). The greater energy intensity of 90 kWh/kg for the *Siemens* hydrochlorination approach, versus 12 kWh/kg for the FBR approach, explains the differences in energy costs that can be seen in the figure. Within the U.S., electricity rates as low as $0.025/kWh (current U.S. dollars) can be had for industrial customers who are able to locate near hydroelectric dams—currently the lowest cost method for generating electricity in this country. Not surprisingly, given the substantial energy requirements, most poly production within the U.S. occurs near these low-cost electricity sources. Accordingly, the energy costs represented in Fig. 3 are calculated on the basis of a $0.025/kWh electricity price assumption. (For all other manufacturing steps in the c-Si module supply chain, which do not necessarily occur in

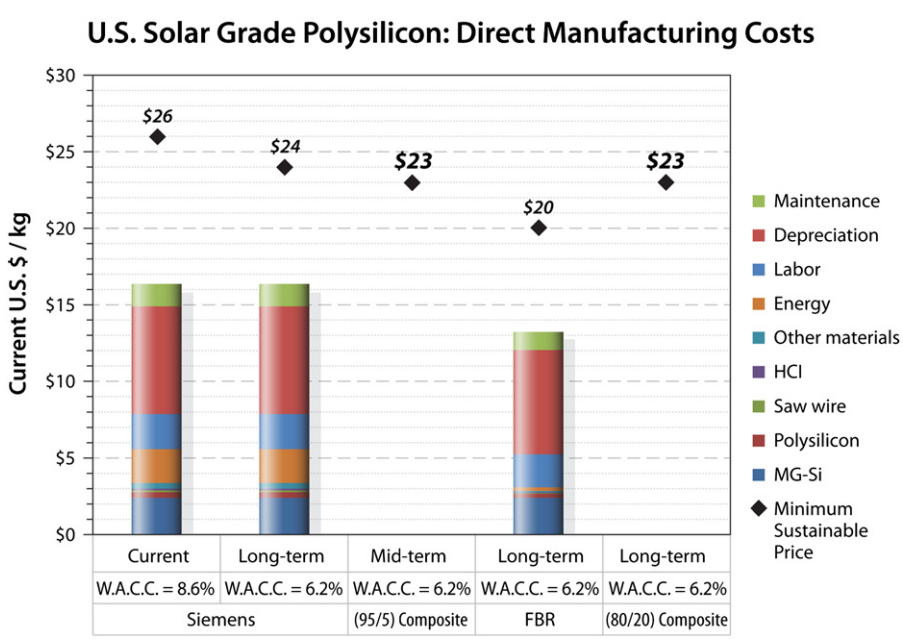

Fig. 3. Model results for polysilicon production costs and minimum sustainable prices, for a U.S.-based 15,000 MTPA production facility with onsite TCS production. The two most commonly employed methods within the U.S. are shown, where the minimum sustainable prices were derived using the specified WACC and *Siemens* chunk/FBR material ratios. The 'polysilicon' and 'saw wire' components correspond to the *Cz* pulling of ingots and the shaping of filaments. For several of the major inputs used for the calculations, please also see the spreadsheet contained within the supplementary information.

such locations, we assume an electricity price equal to the U.S. national industrial average electricity rate of $0.069/kWh [16]).

While some analysts believe that polysilicon producers will be forced to lower long-term contract prices to less than $20/kg by as early as last year [9], by our estimates price demands below this level are not sustainable over the long-term. An increased use of the FBR material may appear to provide the most likely pathway to getting there; however, new FBR facilities are not being built at a rate that is commensurate with a wholesale replacement of polysilicon chunk facilities. And so, despite the ostensible cost advantages in making it, the extent to which the FBR material will contribute to global polysilicon supplies appears to be limited for at least the foreseeable future. Per projection from one large polysilicon producer, we assume that its contribution to global supplies will likely be limited to just 2% in the short-term, 5% in the mid-term, and 20% over the long-term. After incorporating these estimated FBR contributions, and after deriving the minimum sustainable prices needed to meet the mature market WACC, we estimate the long-term minimum sustainable price for U.S. based polysilicon production to be around $23/kg (Current U.S.), with an 80/20 mix of *Siemens* chunk/FBR granules.

4. The Czochralski process for pulling monocrystalline silicon boules, followed by cropping and wafering

4.1. Technical overview

The next step in the c-Si supply chain often takes place in a separate location from polysilicon feedstock production (even in the case of vertically integrated firms), and involves melting polysilicon chunks (and FBR granules, if utilized); forming a *Cz* boule (or 'ingot') from the melt; cropping the crown, tail, and sides of that ingot into a precise shape that minimizes scrap losses; and the cross-sectional sawing of the boule into individual wafers.

This process is begun by first immersing a rotating crystalline silicon seed crystal into molten silicon ($T_{melting}$=1410 °C). The seed serves as a template for the growth of a nearly perfect single crystal, set precisely in length and diameter by the vertical pull rate of the seed from the melt, the amount of polysilicon that can be melted in the crucible, the temperature gradient within the crucible, the rotational velocity of the seed, and the amount of counter-rotation by the crucible [32]. The precise dimensions of the formed monocrystalline boule are carefully calculated on the basis of minimizing material scrap losses in the subsequent cropping and wafering steps, and in consideration of the mechanical fidelity of wafers for all steps through module assembly.

In Fig. 4 a typical process and material flow is shown for producing today's standard wafers having a thickness of 180 μm. For making such wafers, the usual body diameter of an uncropped ingot would be 205 mm, and a representative length would be around 2100 mm (including the tapered ends of the crown and tail). Of course, for the purpose of creating uniform wafer sizes the crown and tail of the ingot cannot be used; they must be cropped off before the wafering process can begin. After cropping a 2100 mm long ingot with a band saw, the final length would be about 1700 mm (with 200 mm equally cut off from both the crown and tail). By sawing off chords of material down the length of the boule, the cylindrical shape is then sawn into that of an elongated square brick with rounded corners: a so-called 'pseudo-square' shape. After accounting for the corner losses, the total cross-sectional area of the shaped brick (and, correspondingly, the pseudo-square shaped wafers that are later used in cell and module assembly) is most commonly set to a standard 237 cm^2, with a flat-edge width equal to 156 mm.

The bulk scrap that is generated during the cuts of the crown and tail, and the chord scrap, is readily remelted during later rounds of ingot casting (after being broken into chunks and chemically washed of the native oxide that develops while the chunks are exposed to air). But, in both the cutting of the bulk scrap and in the sawing of wafers, the kerf sawdust that is also generated is generally unusable when the incumbent approach of standard-wire cutting is employed. With the typical wire diameter of 120 μm producing 130 μm of kerf loss for each cut, and with the large number of cuts undertaken to produce all of the wafers that can be taken from a boule, this kerf loss continues to remain the most significant contribution to the final net material loss in wafering [33].

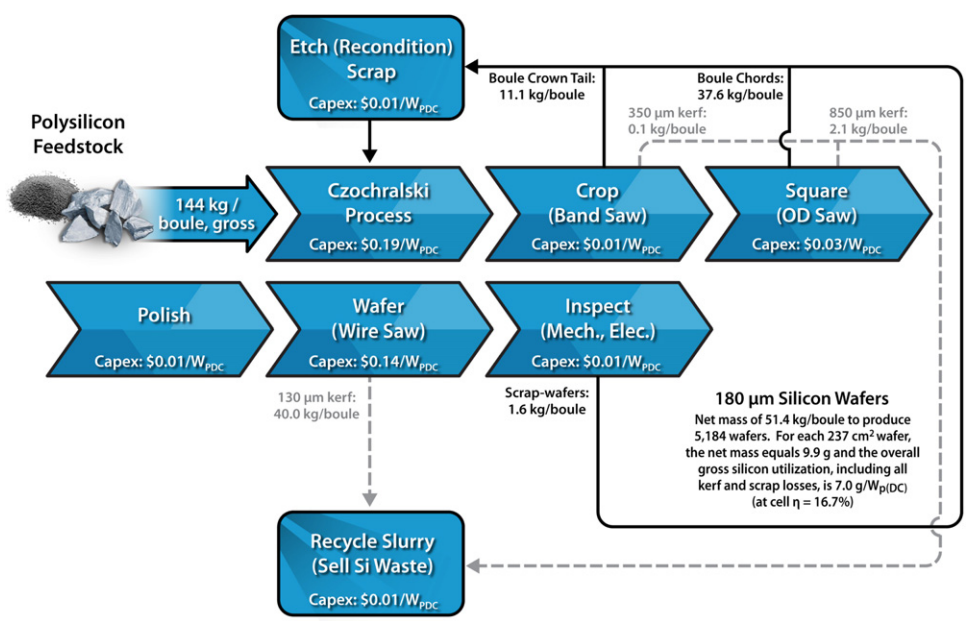

Fig. 4. Process and materials flow for standard *Cz* growth of monocrystalline silicon ingots and subsequent cropping, squaring, and wafering. Typical material losses in production are shown on the outside of the processing steps, where the solid scrap generated through sawing of the boule crown, tail, and chords is recycled for further ingot pulls; but the kerf loss in sawing is not. The given 'Capex' numbers within each step refer to the associated capital equipment expenses divided by the annual production capacity of the facility, with an assumed solar cell power conversion efficiency of 16.7%.

Because this kerf material loss can translate to a significant cost penalty (particularly when polysilicon prices are high), there remains a strong desire to eliminate it. Since no wire-sawing process can completely eliminate at least some amount of sawdust; perhaps it could be just as well to recycle the kerf. For this purpose, one promising trend to have recently surfaced in ingot cropping, squaring and wafering is to switch from the current industry-standard steel cutting wires to cutting wires that have industrial-grade diamond particles attached to them. By using such wires, the diamonds can serve as the abrasive elements for the cutting, thereby replacing the suspension of SiC particles that serve as the abrasive elements in the incumbent standard-wire cutting approach. This is advantageous in terms of the cutting fluid chemistry: Because the SiC particles used in standard cutting must typically be distributed as a slurry with polyethylene glycol, the diamond-wire approach offers a more likely pathway to kerf recycling because the cutting fluid can instead be a simple aqueous surfactant solution—thus yielding a kerf material that possesses a greatly reduced chemical contamination [34–36]. The benefits of diamond-wire sawing do not stop there, however, as it also offers a longer wire life, a faster cutting rate and lower cost for the cutting fluid. In the next section we quantify how these benefits might lead to a reduction in total wafer manufacturing costs.

4.2. Cost analyses of ingot casting and wafering

In Fig. 5, we present wafer production costs for the first half of 2012 (leftmost bar), and the calculated future costs for producing, cropping, shaping, and wafering standard monocrystalline silicon ingots.

The labor costs for wafer production are calculated to be substantial. This is because so many of the steps in wafering are currently difficult to automate. First, the cropped and shaped bricks are manually glued to a glass substrate before being placed into the wafering machine. After the brick is cut, the wafers are released from the glass by immersing the entire unit in an aqueous solution formulated to dissolve the glue. The result is a stack of thin wafers that adhere to each other by virtue of the solution's surface tension. To separate the wet wafers from one another, without incurring high mechanical yield losses due to wafer breakage, requires a level of dexterity that has so far been most easily achieved not by robots but by human hands. And so, with relatively high labor rates, the labor cost for the wafer singulation step in particular is calculated to be nontrivial for a U.S.-based wafer manufacturing facility.

From the first half 2012 case represented in the leftmost bar in Fig. 5, if the technical improvement opportunities described in the last section could be successfully demonstrated in commercial production then there could be a trend towards lower wafer manufacturing costs. In addition to implementing these technologies in wafer production, a long-term reduction in the price for polysilicon—to something more like its minimum sustainable price—would also provide a very obvious benefit toward reducing the costs for making wafers.

Over the course of 2011, the global monthly average spot price for polysilicon varied between $80/kg and $30/kg. Meanwhile, typical contract prices were around $50/kg [8]. After an aggressive

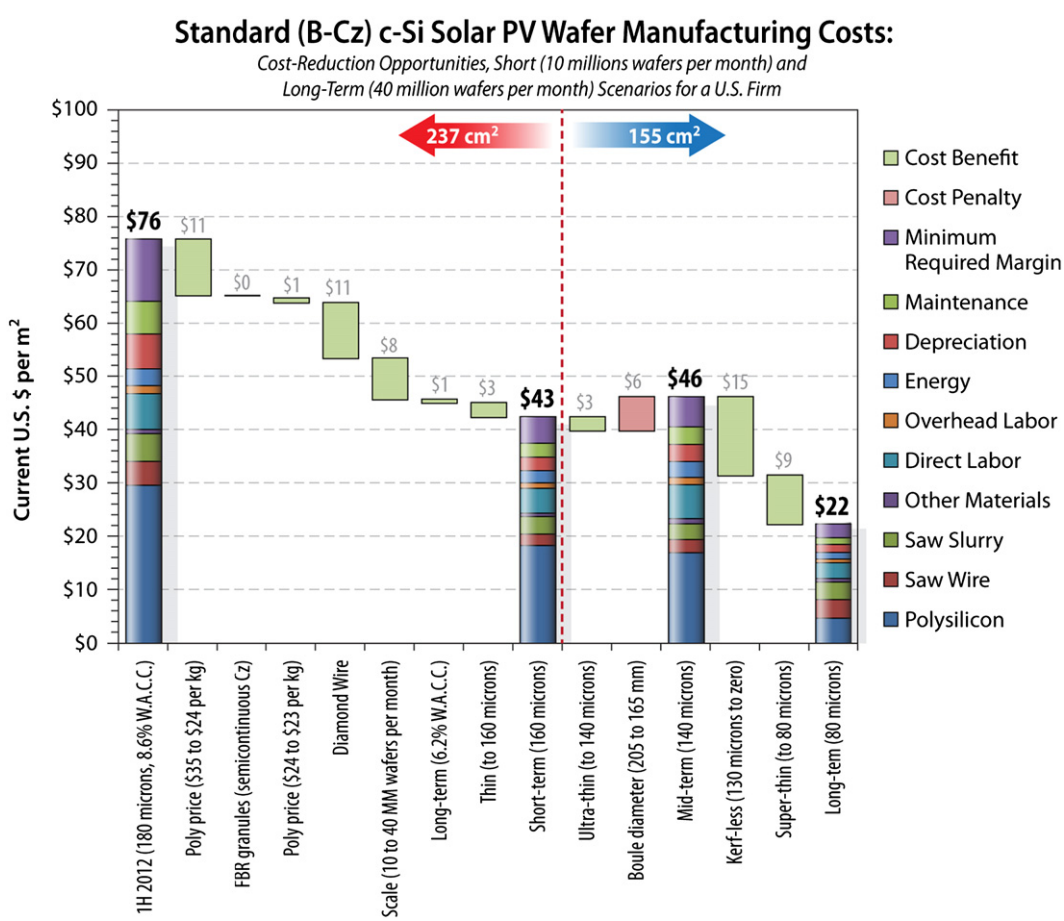

Fig. 5. Estimated current cost and minimum sustainable price (in $/m^2) for producing standard wafers via the *Cz* pulling of silicon ingots and subsequent cropping and wafering (leftmost bar). The waterfall chart then quantifies the specific cost reduction opportunities–and a cost penalty–for each implemented technology described in the text. The modeled facility size is set for an annual production level of 120 million wafers per year in the 1 H 2012 case, and 480 million wafers per year in the long-term case.

worldwide build-out of polysilicon plants, in 2012 poly suppliers had to offer significantly lower prices. At the beginning of the year, typical expectations for spot prices were around $27/kg and typical contract prices were around $41/kg [37]. At the 50%/50% spot/contract prices that wafer manufactures typically work with, at the start of 2012 the blended price for polysilicon would then have been around $34/kg for global wafer production—still noticeably higher than the $24 and $23/kg minimum sustainable polysilicon prices derived in Fig. 3. If the final blended polysilicon price to wafer manufacturers could drop to those minimum sustainable levels, relative to a $34/kg blended poly price there could be a total reduction in materials cost that would work out to around $12/m^2 of wafers produced. As foreshadowed in the Introduction, to obtain such minimum sustainable prices in long-term blended supply contracts may require something like a vertical integration strategy—in addition to globally stable poly prices.

Another pathway to wafer cost reductions could be through an increased use of FBR polysilicon granules. While it is currently not possible to rely exclusively upon FBR granules in the pulling of *Cz* ingots—because current equipment configurations require that the larger *Siemens*-based chunks be present in the initial material loading and melting steps—their increased use in this step of the supply chain enables the effective uptime of the capital equipment to be improved through semi-continuous feeding. This then enables a decline in the overall depreciation expense because the effective 'uptime' of the equipment is improved. However, the benefit is not calculated to be that much. Relative to the benchmark case of just one daughter ingot per initial charge if solely using polysilicon chunk, if three daughter ingots per initial charge could be achieved then the associated savings could be around $0.35/m^2 of wafers produced.

In Table 2 we highlight the major cost-of-ownership considerations for both today's standard wire approach as well as diamond-wire sawing. Although the cost-per-meter of the diamond wire is higher than the standard cutting wire, by utilizing the diamond-wire approach there may be substantial benefits in the final net cost for producing wafers. Before consideration of the additional potential benefit of kerf recycling, the diamond-wire approach may enable an overall savings in wafer manufacturing of around 15% (if the cost of diamond wire can fall to within twice the cost of standard wire, and assuming equal capital equipment costs). If diamond-wire sawing also leads to kerf recycling, an additional cost reduction of up to $15/m^2 may also become possible, depending upon the final costs for the recycling process. (Because the costs for kerf recycling are not currently known, within the Supplementary Information we provide a curve showing the sensitivity of wafer costs to the costs for kerf recycling).

The contribution of future wafer costs to total future module manufacturing costs is generally expected to be lower because the silicon utilization—that is, the grams of silicon needed per Watt of solar cells produced—is expected to improve as cell efficiencies rise and as the wafer thickness is reduced. As cell efficiencies greater than 20% have already been demonstrated on sub-50 μm substrates [38–40], it may even seem that this could happen at any time. But beginning with wafering—and also continuing into cell and module assembly—moving down from today's standard wafer thickness of 180-μm to something much thinner is quite challenging for currently available sawing and handling equipment. The principal challenge is that, for wafering and for all downstream steps, when using today's manufacturing equipment the mechanical yield losses generally increase as the wafer thickness is reduced—and can ultimately become quite intractable for wafer thicknesses below about 80 μm [41].

To partially address this yield loss challenge, the area of the wafer can be reduced while the thickness is reduced. The current guideline is that once the thickness of a wire-sawn wafer is reduced to 140 μm, the boule diameter and final cross-sectional area of a wire-sawn ingot should be reduced from a standard 205 mm (237 cm^2) to 165 mm (155 cm^2). While the guidelines for boule diameter for wafers having a thickness between 80 and 140-μm are less well known, under the best of circumstances it may be possible to make wafers having the same 155 cm^2 area. It may also be possible to achieve a condition of mass balance in that the length of the ingot can be increased, as the diameter is reduced, in order to fully utilize the material capacity of the *Cz* crucible. Additionally, the vertical pull rate can also be slightly increased when making a smaller diameter boule [42]. (From the data provided by a relevant equipment supplier, the achievable vertical pull rate of the boule body is calculated to be around 43 mm/h for 237 cm^2 wafers, and around 49 mm/h for 155 cm^2 wafers). Even so, there is an overall net cost penalty associated with the smaller area wafers because there is an overall lower throughput in the *Cz* pullers and in the sawing equipment. After accounting for these factors, with the relevant collaborator-provided inputs this small-area penalty is calculated to be around $6/m^2.

There are, of course, several other ideas beyond what we have outlined that also seek to improve the net silicon utilization. These include ideas such as the epitaxial growth and lift-off of film silicon, cast wafers, exfoliated wafers, ion-based cleaving approaches for wafer separation from an ingot, etc. [43]. Each of these approaches would have its own proprietary cost structure for wafer production, in addition to other considerations for how such wafers would integrate into the overall c-Si supply chain. Knowing that there are other approaches to wafer production that are being considered, it is worth noting to the reader that the wafering cost estimates that have been presented are only for the technology cases that have been described.

Table 2
Cost of ownership assumptions for the standard and diamond-wire saw approaches to cropping and wafering monocrystalline silicon boules.

Cost of ownership input assumptions for the standard-wire and diamond-wire approaches to cropping, squaring and wafering monocrystalline silicon ingots

	Standard wire	Diamond-wire
Wire diameter	120 μm	120 μm
Kerf loss per cut	130 μm	130 μm
Cutting rate (mm/min)	0.37	1.1
Cutting fluid and cost	SiC in PEG $1.40/kg - $2.00/kg	Water with surfactant ($ 0.39/1000 l)
Wire cost	$2.80/km	$5.60/km
Wire life (cm^2 of wafers produced per meter of wire)	24	80

5. Cells

5.1. The standard monocrystalline silicon solar cell

The majority of c-Si solar cell production is currently based upon a very standardized process that is intended to make a *p*-/*n*-electrical junction on the entire front surface of the wafer and a full-area aluminum-based metallization on the back [44]. A representative series of steps for making such cells is shown in Fig. 6.

First, because wafers are typically received from multiple supply sources, and because they can be damaged during sawing and shipping, all incoming wafers should be tested to ensure that they would provide a foundation for acceptable cell efficiencies. Specifically, for the purposes of wafer metrology, measurements of the minority charge-carrier lifetimes would be quite informative [45]. But due to an anticipated slowdown in manufacturing throughput if every single wafer was to be tested—primarily

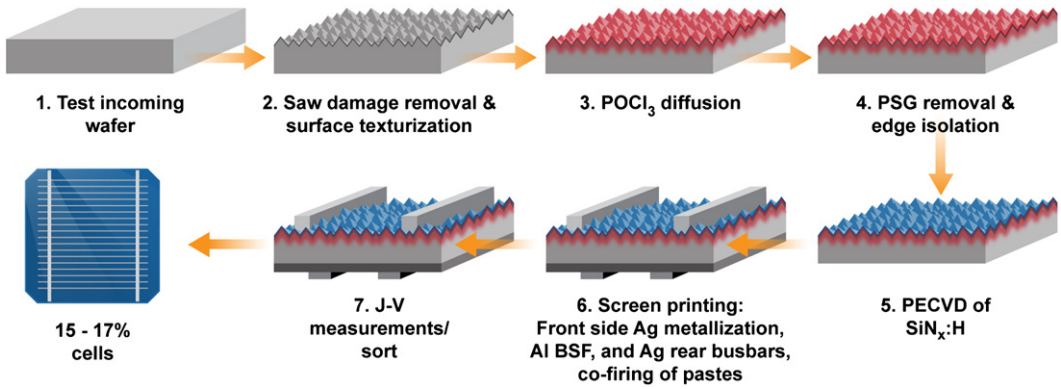

Fig. 6. Process flow for fabricating a standard c-Si solar cell.

because it can be difficult and time-consuming to decouple the effects of surface recombination from bulk recombination when interpreting such measurements—it has so far proven elusive to cost-effectively standardize the in-line metrology of all as-cut, unpassivated wafers immediately after they are taken out of the box [46,47]. While this is rapidly changing, particularly in the more state-of-the-art cell manufacturing facilities, in the current standard process as-received wafers are more simply and quickly screened for their dimensions and for any physical damage. Most commonly, the standard process typically employs rapid throughput wafer-imaging systems that detect the defect states associated with microcracks (by mapping attenuations in IR transmission), and reject any wafers containing microcracks that are too large.

Because the wire sawing process is extremely abrasive, if left untreated an as-received c-Si wafer would retain a high density of unpassivated surface defects that would actively facilitate electron–hole recombination in an illuminated cell. To ameliorate this problem, and to partially remove the microcracks at the surface that might compromise the wafer's resilience to breakage during handling in cell processing [35], as the second step in the standard process a wet bench chemical treatment is utilized to etch away between 5 and 15 μm of saw damage from the top surface of the wafer. This is typically achieved by exposing the wafers to an aqueous solution of NaOH or KOH with IPA. With the alkali metal only being a spectator ion, the etching reaction proceeds as follows [48]:

$$Si + 2H_2O + OH^- \rightarrow HSiO_3^-(aq) + 2H_2$$

The etch rate of this chemical reaction is different for different crystallographic orientations. Due to these anisotropic differences in etch rates, the originally flat wafer surface is etched into a morphology of pyramids having a random distribution in size. Fortuitously, these pyramids can provide the foundation for front-surface light trapping [49]. At the conclusion of this wet bench chemical processing step the surface is then ready for the formation of the topside *p*-/*n*- electrical junction.

In a well designed solar cell, the surface and bulk electric fields should work to usher all of the photogenerated electrons and holes towards their appropriate electrical contacts. This is most generally achieved by establishing a *p*-/*n*- junction across the region that absorbs the most light. The standard c-Si solar cell is generally made with a boron doped (*p*-type) base wafer. The formation of the *n*-doped region—the so-called the 'emitter' region—is formed over the entire topside of the wafer as the doping characteristics are inverted from *p*- type into *n*- type by the high-temperature drive-in of phosphorus [50]. In this step, the wafers are exposed to phosphorus oxychloride ($POCl_3$) gas within a quartz tube furnace and then heated to a high temperature, typically between 800 and 900 °C, in order to activate the diffusion of phosphorous into the wafer.

During this diffusion step, the surface of a $POCl_3$ treated wafer becomes glassy. Because this amorphous layer makes it difficult to make a good electrical contact to the bulk silicon, and because its properties change after being exposed to moisture, it is generally necessary to include another processing step for the removal of this thin phosphosilicate glass (PSG) layer. For this purpose, an HF dip is typically used. So that only one *p*-/*n*-junction is formed at the top of the solar cell, another treatment commonly called 'edge isolation' is also typically carried out. The purpose of the edge isolation treatment is to remove the shallow phosphorus diffusion that creeps onto the wafer edges and backside, which is unavoidable even though it was only the topside that was primarily exposed within the tube furnace [51]. In the wet bench chemical approach to edge isolation, the same manufacturing tool can be used as for PSG removal.

While the random pyramid surface texturing is helpful, depositing an additional layer that possesses a sufficiently different index of refraction from silicon can help to reduce the reflection of light even further. As a standard material, hydrogenated silicon nitride (SiN_x:H) is able to serve as such an antireflection (AR) coating. The plasma-enhanced chemical vapor deposition (PECVD) approach is currently the most widely employed method for depositing this material. In this process, the AR coating is formed during the plasma-activated reaction between silane (SiH_4) and ammonia (NH_3) gases that are introduced into the reactor chamber.

Because it has generally been more cost-effective than vacuum-based metallization approaches like evaporation or sputtering, the screen printing of Ag and Al pastes for the formation of the front and rear electrical contacts has been in use by the c-Si industrial community since the 1970s [52]. In this process a conveyer belt moves c-Si wafers along a queue where they are picked up, either by a robotic or human arm, and placed onto a printing table. An H-pattern screen that is mounted in an aluminum frame is then overlain on the frontside of the cell and the metallization paste is squeegeed over the wafer surface with a defined pressure. In today's screen printers this handling and printing process can be repeated at an impressive net rate of 1–2 s per wafer, including time for wafer placement and removal. After the front-side screen-printing, the wafers are moved into a low-temperature (≈ 200 °C) drying oven, and the wafer is then moved on to another table for a three-step printing sequence of the rear side Al paste and Ag/Al rear busbars. At the end of the printing steps, in order to drive off the undesired additives used to make the metal pastes, the entire cell assembly is typically co-fired at around 810 °C [53]. At this temperature, lead borosilicate glass frit ($PbO–B_2O_3–SiO_2$) contained within the Ag paste etches through the SiN_x:H layer to form a direct bond and electrical contact with the underlying emitter region.

As the final step in the standard c-Si solar cell processing sequence, the current–voltage (J–V) characteristics are measured for each cell that is produced on the line. In order to minimize current mismatch losses between cells when they are series-connected into modules, they are binned according to their current density at maximum power point.

5.2. Introduction of potential pathways to improve efficiencies beyond the standard c-Si cell

While the standard approach to cell processing has been the dominant manufacturing strategy for quite some time, it is increasingly clear that it will become necessary to lower costs even further, in order to remain competitive within the future landscape of PV. For all steps within the c-Si supply chain as well as at the installed systems level, there is little choice but to call upon gains in efficiency in order to achieve these ends—and it appears that the standard cell processing approach will ultimately not be able to deliver the 20–25% power conversion efficiencies that other industrially-relevant manufacturing processes are capable of delivering [54].

To define this important consideration, the efficiency of a solar cell is most generally calculated as follows:

$$\eta = \frac{\text{MaximumCellPowerOut}}{\text{PowerIn}} = \frac{FF \times J_{SC} \times V_{OC}}{\int_{\lambda_{min}}^{\lambda_{max}} \text{E}_{\text{AM1.5}}(\lambda)d\lambda} = \frac{FF \times J_{SC} \times V_{OC}}{1000\text{W/m}^2} \tag{9}$$

where AM 1.5 represents a modeled profile for the number of photons expected for each wavelength within the solar spectrum after it passes through the earth's atmosphere. By carefully specifying the atmospheric conditions in an internationally standardized way, which falls within the expectations of sunlight received at mid latitude on a clear day, the total integrated energy content of this modeled profile can be set equal to 1 kW/m^2. The FF, J_{sc}, and V_{oc} represent the respective efficiency parameters of fill factor (unitless), short-circuit current density (in A/m^2, or mA/cm^2), and open-circuit voltage (in Volts). The product of these three gives the maximum power under standard test conditions. For c-Si, in consideration of the absorption profile of the semiconductor in comparison to the AM 1.5 spectrum, as well as the factors that limit the V_{oc} and FF, the full efficiency potential under standard test conditions is around 29% for a 100-μm wafer [55].

If it assumed that the solar cell follows ideal diode behavior with superposition, the V_{oc} parameter can be analyzed after knowing the inputs for Eq. (10), where n, k, and q are constants at a fixed temperature [56]:

$$V_{OC} \approx \frac{nkT}{q} \ln\left(\frac{J_{SC}}{J_O}\right) \tag{10}$$

As Eq. (10) shows, maximizing the V_{oc} implies maximizing the J_{sc}, while at the same time also minimizing the J_0.

This J_0 parameter—most commonly called the saturation, or recombination, current density—is a broad representation of the overall net rate of electron-hole recombination within a solar cell. While the theoretical limits to the J_0 and V_{oc} in a c-Si absorber are estimated to be around 0.27 fA/cm^2 and 0.845 V [54], the measured values for a standard cell are noticeably different from these. There is a long list of causes: radiative, conduction-band to valence-band recombination (i.e., emission); defect-mediated recombination on the front, edge, and back surfaces (which is but one form for Shockley–Read–Hall, or SRH, recombination); SRH recombination from states generated by bulk defects; metal-to-silicon contact recombination; p-/n- electrical junction (or 'depletion layer') recombination; Auger recombination; and the list goes on. In the end, the lower limit to the J_0 is set by the radiative recombination term because, at the thermodynamic limit of a perfect detailed balance, the rate of electron–hole pairs being put back into the solar cell must equal the rate that is generated by light absorption [55]. (For a more complete description of this lower limit, and for a history of c-Si technology development, please see Ref. [57].)

As for the other, non-radiative, sources for recombination, there are a multitude of potential remedies that are commonly recognized. We introduce several of them in Table 3.

By calling upon the same underlying assumptions used to derive the open-circuit voltage expression, it is also possible to simplify the discussion of the short-circuit current to the following [56]:

$$J_{SC} = \int_{\lambda_{\min}}^{\lambda_{\max}} \text{QE}(\lambda) \times \text{AM1.5}(\lambda)\text{d}\lambda \tag{11}$$

where the QE(λ) term, called the external quantum efficiency, is the probability for generating and collecting an electron at the specified wavelength of incident light. In the typical reference case, the maximum possible current (i.e., at short-circuit condition) of the solar cell can then be derived by integrating these probabilities over the entire AM 1.5 spectrum. As it is for the V_{oc}, there are architecture-specific origins for the observed losses in the J_{sc} of a standard c-Si solar cell. Several ideas for improving them are captured in Table 3.

In arranging the table some of the major technical improvement opportunities that are known are organized into three general technology groups. The first difference between these technology groups is in the choice of base doping within the wafer: p-type for Group 1, and n-type for Groups 2 and 3. The hypothetical cells for each technology group capture several of the currently known opportunities for creating more efficient c-Si devices and, in principle, could be made with equipment that is currently available for industrial-scale manufacturing. Nonetheless, the model cell architectures, the underlying efficiency assumptions, and following processing flows associated with each have not necessarily been commercially demonstrated in their entirety, and were explicitly designed for cost-modeling purposes only.

5.3. Technology group 1: front-side metallization on a p-type Cz wafer (20–22% cell efficiency)

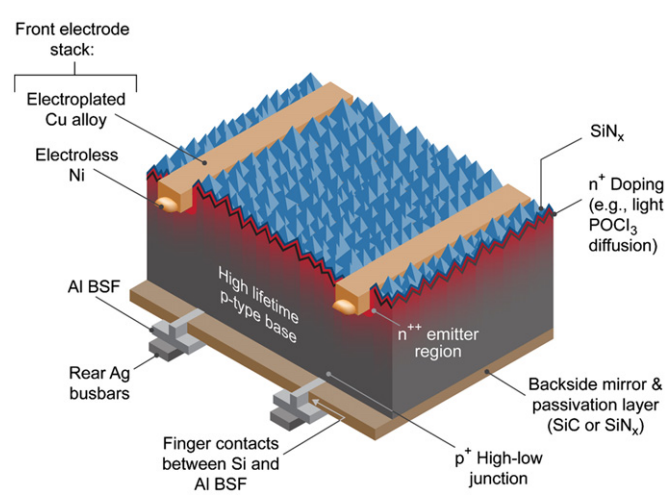

5.3.1. Front metal contact buried into the wafer with a locally diffused emitter

The highest efficiency c-Si solar cell to date, at 25% [93], is based upon an architecture called the Passivated Emitter Rear Locally-diffused (PERL) cell. The record efficiency mark for this cell has been in place since 1999, and, although it is cost-prohibitive to precisely replicate all aspects of the PERL cell, several of its underlying concepts

Table 3
Overview of several technical improvement opportunities, organized by technology groups, that are available to improve the efficiency of c-Si cells and modules. The assumed cell-to-module derate is 89% for the calculation of the module efficiencies shown in parentheses at the bottom of the table, although this value may be improved with changes to the assumed standard module design.

Cell performance parameters	2011 Standard cell (*p*-type base)	Technology Group 1 (*p*-type base)	Technology Group 2 (*n*-type base)	Technology Group 3 (*n*-type base)
Short-circuit current density: J_{SC} (mA/cm^2)	**35**	**38** • Backside optical mirror [58] • Higher aspect ratio front gridlines [59] • Buried front metal contacts[60] • Selectively diffused emitter junctions [61]	**41** • Reduce front-side shadowing losses by moving contacts to the back [62] • Improved light trapping through novel surface texturing and higher internal light reflection [63,64] • Lightly doped FSF [65] • SiO_2 passivation[65]	**40** • Develop a TCO with reduced free-carrier absorption [54,66] • Develop a heterojunction window layer with reduced absorption [66]
Open-circuit voltage: V_{OC} (V/cell)	**0.62**	**0.70** [67] • Selectively diffused emitter junctions [61] • Improve wafer quality: alternative dopants or magnetic Cz [69–71] • Improve surface and bulk passivation [62,72]	**0.74** • Ion implantation for precise control of dopant profiles [73–75] • Use tightly focused metal-to-Si contacts in order to reduce contact recombination losses [76,77] • Use *n*-type wafers with minority carrier lifetimes approaching 10 ms [78] • Improve back, front, and edge surface passivation [55,79,80]	**0.75** [68] • Use *n*-type wafers with ms minority carrier lifetimes [78] • a-Si:H/c-Si heterojunction surface passivation [81]
fill factor: *FF* (%)	**78**	**80** • Improve conductivity (σ) through electroplating [82] • Develop and improve new metal and selective emitter paste chemistries [83,84] • Selectively diffused emitter junctions [61]	**82** • Reduce resistive (I^2R) losses, without compromising optical losses, by covering more solar cell area in a back-contact scheme [54]	**80** • Use *n*-type wafers with ms minority carrier lifetimes [78] • a-Si heterojunction surface passivation [54] • Use a TCO for charge-carrier transport and anti-reflection coating, and develop a new one with a higher electrical conductivity [85]
AM 1.5 power conversion efficiency (%):	**17% cells (14.5% modules)**	**20–22%** [86–89] **(18.7%)**	**25%** [90] **(22.4%)**	**24%** [91,92] **(21.4%)**

are clearly appearing within many of the new equipment designs and industrial research and development programs [87,89,94]. Because the cell is also based upon a boron doped (*p*-type) base, making it quite amenable to the standard industrial cell processing approaches already in place, we have included an industrially scalable derivative of the PERL architecture as part of our c-Si roadmap.

While not the only contributor to the high cell efficiency, the PERL concept incorporates the idea of a heavily doped emitter region that is narrowly focused at the point of contact between Si and the frontside metal, in addition to a lightly doped region over the entire wafer front surface [95]. Today this design is more commonly called a 'selective', rather than a 'locally diffused', emitter; and it possesses several advantages over the standard cell architecture—primarily in optimizing the electrical connection between the frontside metal and silicon without also creating unnecessarily high rates of recombination over the unmetallized regions of the wafer's front surface [50,96]. This selective emitter profile is achieved by creating two different doping densities within the cell: *n*-type doping (with carrier concentrations, N_d, on the order of $10^{19}\,\mathrm{cm}^{-3}$) over the entire wafer front surface, and n^+ doping ($N_d \approx 10^{20}\,\mathrm{cm}^{-3}$) directly at the line of contact between the metal and the narrow emitter region. In comparison to the standard emitter profile, this selective emitter doping profile makes it possible to establish an electric field within the device that more efficiently ushers photogenerated electrons and holes toward their appropriate electrodes—while also greatly reducing the probability for recombination as the charge carriers move between the silicon and the metal electrodes [97]. The overall expected result when employing such a design is a lowered value in the overall J_0 (and, therefore, a correspondingly higher V_{oc}), as well as a slight benefit to the J_{sc} of the cell due to a higher quantum efficiency of blue photons [61,98].

There are numerous manufacturing processes currently under development that can deliver cells of this type. These include the industrially-relevant options of either screen-printing dopant pastes [84,98], or using the laser-assisted doping of a wafer from a stream of H_3PO_4 [99,100], to form the n^+ region. As for how they might be implemented, we have incorporated either nominally cost-equivalent option as a step within a process flow designed around an industrially scalable derivative of the PERL cell (Fig. 7). While not used within our specific process flow, there are also other options for making selective emitter contacts, including a heavy doping of the entire wafer frontside (followed by selective etch-back of the cell's surface in all regions except where the metal contacts are to be printed), or screen-printing a heavily doped silicon nanoparticle ink [101].

In creating the process flow for Technology Group 1, it is assumed that standard cell processing steps 1–5 could be retained. That is, standard cell processing steps 1–4 could be called upon to establish light *n* doping of the wafer front surface; and step 5 could be called upon for incorporation of SiN_x:H (because, as in the standard cell, this material is beneficial as an anti-reflection coating and for mitigating J_o due to front-surface recombination [54,102]).

After the PECVD step, it could then be appropriate to form a network of trenches in the silicon wafer with a laser ablation step. In this process, a focused laser that is powerful enough to create localized heating above the melting temperature of both silicon ($T_{melting}$=1410 °C) and silicon nitride ($T_{melting}$=1900 °C) is used to create a pattern of trenches having a precisely set depth and width [99,103]. Such an ablation process may create a thin layer of damage on the walls of the groove, however, which should be removed with something like a NaOH etch. (Otherwise, if left untreated, dislocations in the local crystal structure that are generated by the lasers can potentially glide into the bulk during subsequent thermal processing steps.) We have included this laser damage removal step as step 7.

The laser groove is envisioned to be advantageous for either the printed dopant paste or aqueous-based approaches to emitter drive-in, primarily because it also establishes a pathway to higher aspect ratio grid lines on the front of the cell—a benefit we discuss next. It is worth noting that steps 6 and 8 would likely require separate laser stations.

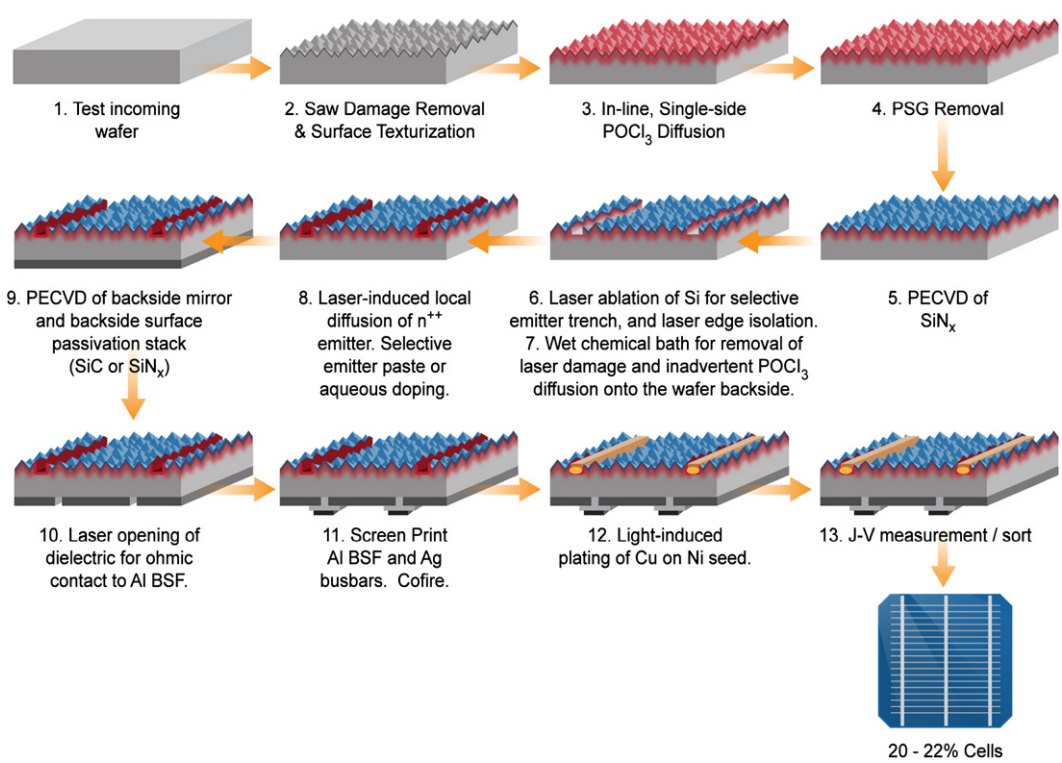

Fig. 7. Model process flow for fabrication of Technology Group 1 cells.

5.3.2. Moving from screen-printing to electroplating for front-side metallization

For the most part, there appears to be little rush to replace the standard screen-printing of Ag. It is, after all, a commercially proven, high-throughput process with good alignment control and low wafer breakage rates. Just as importantly, it also provides a metal contact having reliable adhesion to the solar cell over the entire lifetime of the module. Yet there are multiple reasons as to why this approach to frontside metallization is ultimately unsustainable.

One commonly expressed reason is that, at some point, the c-Si technology may eventually become a victim of its own success, in that the demand for Ag from it alone may drive its price to an unacceptably high level. Without confirming or denying that particular line of conjecture, the question does need to consider that there will more than likely be a noticeable decline in the required grams-of-Ag-per-watt in c-Si manufacturing, because there are other approaches to metallization—beyond just simple screen-printing—that can replace a lot of the Ag with much cheaper metals.

The next step of our model process flow incorporates the light-induced plating approach for depositing a copper-alloy grid on top of an electroless seed of nickel [60,104,105]. This metal layer stack could conceivably move c-Si cells away from the full screen-printing of silver and toward the more scalable alternative of a primarily Cu-based metallization. Even if this choice is developed, as it already has been within the integrated circuits industry, we have included the Ni layer to be but one example of a material that will likely be needed for adhesion, to prevent the diffusion of Cu into the cell (which is detrimental to the reliability [106]), and to also serve as an electrical channel for the deposition of a Cu alloy from solution.

In the first step of this method, a very narrow seed and barrier layer around 5 μm in width (much more narrow than today's usual full Ag line widths of 60–100 μm), and around 1.0 μm in height would be deposited in the trenches via electroless plating, screen-printing, ink jet printing, or aerosol printing [59,82]. (Our cost calculations for this step of the process flow are based upon an electroless plating of the barrier layer, with SiN_x:H serving as a built-in dielectric plating mask.) The light induced electroplating process could then conceivably be used to thicken the line, either with the same metal, or, if need be, with a different metal (such as a Cu-alloy). This elegant plating process is driven by employing the photovoltaic effect for a cell immersed in an electroplating bath: after applying the necessary reduction potential and exposing the cell to light, the photogenerated electrons flow to the surface and then into metal ions within the surrounding solution [86,87]. The plating process continues until the desired line conductivity (cross-sectional area) is reached, with current cycle times being just a few minutes per cell [67,107]. As to what metals could be used within the alloy, per the model provided by a relevant equipment supplier, our specific cost-of-ownership estimates for such a plating process are based upon a three bath sequence that produces layer thicknesses of 3.0 μm for Cu, 3.0 μm for Sn, and 8.0 μm for Ag.

The light-induced plating approach has been used for quite some time in several higher efficiency cell designs because it has demonstrated an absolute efficiency improvement of at least 0.3–0.5% over screen printing [108]. One of the main contributors to the efficiency improvement has been that this approach to metallization can produce much narrower completed line widths of around 30–50 μm [82,87,108]. This is advantageous in reducing losses to the J_{sc} due to shadowing and reflection of light from the frontside metal. But using narrower metallization lines does require using either a higher density of grid lines on the wafer surface, or lines with a higher aspect ratio, in order to limit resistive losses and to move the same amount of photocurrent as the wider screen-printed lines. It is on this point that the laser groove serves yet another purpose. Controlled by the trench depth and width, the height: width metallization aspect ratio can be moved from around 1:4 to 1:2, and so the line conductivity can be improved while also reducing the total amount of dead area on the cell [86,87]. As final notes in addressing the replacement of screen-printing with a laser-buried groove in the Technology Group 1 cells, this concept has also been shown to reduce the area-dependent contact recombination losses between metals and silicon [82]. Moving away from the dielectric glass frit that is contained within Ag paste should also lead to lower overall series resistance losses within the grid array [83].

5.3.3. Improved front- and back-side surface passivation, back surface field, and backside mirror

With little time to lose, immediately after a solar cell absorbs light it is necessary to drive the photogenerated electrons and holes towards their appropriate electrode terminals before they recombine [109,110]. The average time that is available to move a given charge carrier—before a recombination process would be expected to occur—is called the charge carrier lifetime (τ). The corresponding average distance that the charge carrier will be able to travel during that time is called the diffusion length (L_d). L_d is related to the lifetime by $(D\tau)^{1/2}$, where D, the diffusion coefficient, is around 30 cm^2/s for a minority carrier electron within the bulk of a p-type c-Si wafer [111], and around 11 cm^2/s for a minority carrier hole within the bulk of an n-type c-Si wafer [112].

Over the course of their transitory lifetime, mobile charge carriers will spend relatively more time either diffusing through the bulk of the wafer or in navigating the wafer surfaces. The relative amount of time spent at either depends upon the bulk diffusion length and the wafer thickness and size (in essence, for a given diffusion length, the thinner and smaller in area the wafer is, the more relative time a charge carrier will spend near the front, back, and edge surfaces). While they are on the surface, the mobile charge carriers must traverse a thicket of dangling bonds and other defects that can capture them and trap them in place—at least until they either escape the trap and move onto the next one, or until another oppositely charged carrier happens to come along and the process of recombination occurs. The best strategies for keeping charge carriers away and free from these traps are: (i) to create an electrical potential energy barrier (specifically, a space-charge region around the p-/n- electrical junction and a back surface field) in order to repel a chosen charge carrier away from the surface in the first place; (ii) to reduce the rate of surface recombination by reducing the concentration of one charge carrier type at the surface; and (iii) to passivate all dangling bonds and defect states to the fullest extent possible [102].

Whether or not one has been successful in this regard can be inferred from measurements of the surface recombination velocity (SRV), which is parameterized separately for the front and back. There are several built-in features within a standard cell that lower the SRV to a level such that it can operate at least reasonably well. The p-/n- junction on the wafer front surface is the first such feature. By inverting the doping characteristics from p-type into n-type, the process of $POCl_3$ diffusion creates a wafer front surface where the concentration of one charge carrier (electrons) is significantly higher than the other (holes). After the electrons migrate down their concentration gradient—from the surface into the bulk—the resulting vacancies, or 'holes', that are left behind are positively charged, and this creates an internal electric field within the solar cell which works to repel charge carriers of one particular type (in this case, other incoming holes) away from the surface [56]. With a lower surface concentration of one particular charge carrier in the electron–hole recombination pair, the overall rate of surface recombination is thus lowered. On top of this junction, the SiN_x:H anti-reflection coating also helps to significantly lower the front SRV, from $\approx$250,000 cm/s to $\approx$40,000 cm/s, by passivating surface defects and dangling bonds [113].

There is also another built-in surface passivation on the backside of the standard c-Si cell called the high-low junction—named because the Si at the back region is locally doped by Al to a higher, p^+, doping level ($N_{h+} \approx 3 \times 10^{18}\ \mathrm{cm}^{-3}$) in comparison to the bulk of the base wafer ($N_{h+} \approx 1\text{–}2 \times 10^{16}\ \mathrm{cm}^{-3}$) [114]. As in the case of the emitter junction (at least in the sense of the surface and the bulk possessing different carrier concentrations), this high-low scheme establishes another type of 'charge carrier mirror', or 'back surface field', which works to repel minority carriers away from the rear surface of the wafer [56,115]. In the absence of additional layers to passivate surface defects, however, the typical back SRV realized for a full-area Al back surface field (BSF) in industrial production is still very high, at around 1000 cm/s. With optimization, a simple full-area Al BSF might be able to deliver a back surface SRV around 200 cm/s [114,116].

To quote the publication directly: 'the major advantage of the PERL cells is passivation of most of the cell surface areas' [95]. Even before the champion PERL cell result, the first achievements of efficiencies greater than 20% can also largely be attributed to the implementation of more effective passivation techniques [117,118]. In light of this extremely important efficiency improvement parameter, our roadmap and cost models incorporate enhanced backside passivation techniques, beyond the simple full area Al BSF, for a more efficient Technology Group 1 type of cell. As a minimum requirement, the achievement of $>20\%$ cell efficiencies may very well require that the total back SRV be on the order of just 5–10 cm/s. Meanwhile, the front SRV achieved with SiNx:H still appears to be acceptable for the purpose of achieving such cell efficiencies [102].

The necessarily lower back SRV could be achieved by employing the field-effect passivation approach, whereby additional fixed negative charges are brought into place on the wafer back surface. This could be achieved by depositing an additional dielectric layer(s) between the Si and Al BSF [119]. There seem to be numerous approaches, all of them industry-relevant, for achieving such an end. The first option is to simply heat a wafer to high temperature in the presence of oxygen. A thick layer of SiO_2 would then be produced on all wafer surfaces, as it was in the original PERL cell [95]. While not demonstrated at large production volumes, a lower temperature alternative to achieving an SRV similar to SiO_2 is by the atomic layer deposition of Al_2O_3 [120,121]. This is not the only low temperature option, however. Each one capable of being deposited by industrial PECVD tools, there are at least three other dielectric materials that can also achieve similarly low back SRV: SiN_x:H [102,122], undoped (intrinsic) a-Si:H [102], and SiC [123,124]. Another approach for a Technology Group 1 type of cell might be to try some combination of the above, as is done in the industrial passivated emitter rear cell (*i*-PERC) that uses an SiO_2/SiN_x:H stack [72,119].

For the sake of brevity, our costs modeling results and process flow for Technology Group 1 are limited in scope to the choice of a backside dielectric passivation layer deposited by PECVD. This choice seems logical to us because, with regards to SiO_2 being deposited in a high-temperature furnace, there are known material instability problems for standard B-*Cz* wafers under high temperature oxidizing conditions (which the PERL cell did not suffer because it utilized a float zone wafer) [60,125]. Furthermore, the potential need for additional masking steps only seems to complicate matters for a marginal improvement in SRV [102,126]. Finally, *p*-type *Cz* wafers in particular are difficult to passivate uniformly with SiO_2 under industrial processing conditions [127]. As for the ALD approach, it is certainly intriguing and may be a topic for future research.

Among the options available for the PECVD approach, it is not entirely clear at this time which material would be best. Without commenting further on the choice of a-Si (a perfectly viable option, which we address within the Technology Group HIT), SiN_x:H appears to be a good choice because it is a decent surface passivation layer and it can also serve as a backside mirror of light [58]. As yet another option, SiC is also an excellent passivation layer, as well as backside optical reflection layer, that can be doped to assist with the overall cell mechanics of bringing holes to the cell's back surface [124]. Finally, either choice can ostensibly offer similar solutions to the problem of wafer bowing during the Al firing step [86]. Without having to decide between the two any further, they both appear to be essentially equivalent because they essentially utilize the same capital equipment, have similar thicknesses, and can both be made from precursors with similarly low costs.

The final steps of our modeled Technology Group 1 process flow enable ohmic contact between the silicon wafer and Al. For this there are two conceivable options: either open the dielectric with a laser first, and then screen-print the Al; or screen-print the Al directly on the dielectric first, and then use a laser to drive metal melt-through [128–130]. Both options are estimated to be marginally cost-equivalent, and only require that around 1% of the area on the backside of the wafer be contacted by the metal (which is beneficial for making a backside mirror of light, and in preventing wafer bowing). Within our process flow, we have also assumed that, due to the Al–Si eutectic that is formed during co-firing, an additional laser damage removal step will not be needed [115].

5.3.4. On the need for improved material quality within the base wafer

Boron has a long history as the dopant of choice for standard silicon wafers because its high segregation coefficient works to produce ingots having a very reproducible and uniform distribution of dopant atoms along the entire length of the boule [131]. This remains a highly desirable material property for wafer manufacturers who must sell their product with a specified base resistivity, which is principally set by the dopant concentration within the wafer. However, the inherent efficiency limitations for cells made with the standard B-*Cz* wafers have been known for quite some time [132].

The first complication with standard B-*Cz* can be understood from how a dopant atom incorporates itself into the silicon lattice, because each site that must accommodate a dopant atom is a disturbance to the symmetry of the crystal. Even though the concentration of these asymmetries is usually extremely low (on the order of ppm), if the difference in size between the dopant atom and silicon is significant enough, the process of recombination can be accelerated over a network of the asymmetries, or 'misfit dislocations', that propagate throughout the bulk of a wafer [133]. To begin understanding how these misfit dislocations might vary depending upon the choice of dopant, the atomic radius of silicon is 1.18 Å—and this can be compared to that of phosphorus, with an atomic radius of 1.06 Å; gallium, with an atomic radius of 1.22 Å; and boron, with an atomic radius of 0.88 Å [70,134].

The symptoms of increased bulk recombination—be it from misfit dislocations or other factors—can be diagnosed from measurements of minority charge carrier lifetimes. By these measurements, a few general patterns can be gleaned from the published literature: both before and after a typical cell processing sequence, as well as after illumination, Ga-doped *Cz* wafers (Ga-*Cz*) will retain lifetimes on the order of hundreds of microseconds [70,135]; P-doped *Cz* wafers will retain lifetimes on the order of milliseconds under the same stresses [125,133]; but the lifetimes of the traditional B-*Cz* wafers will rapidly decay from the hundreds of microseconds to the tens of microseconds [125,136–138].

A careful reading of the same literature, however, also highlights a second and more significant explanation for the observed changes in carrier lifetimes within B-*Cz* wafers. During the *Cz* pulling of ingots, it is unavoidable that oxygen from the silica-based crucible is released into the melt, where it is strongly attracted to boron. As the melt

solidifies during the pull, these boron–oxygen pairs can remain in place within the crystal lattice. By first observing that the presence of oxygen very much coincides with the light-induced carrier lifetime degradation effect seen in B-*Cz* wafers, and then by carefully monitoring the concentration of both the boron and oxygen species within a defined ensemble of wafers, it can be seen that the lifetime degradation typically seen in B-*Cz* wafers is due more to the presence of oxygen than it is due to the presence of boron [125,139]. To quantify the experimental observations, the concentration of the metastable boron–oxygen complexes that lead to the lifetime degradation depends quadratically on the concentration of oxygen and linearly on the concentration of boron [140].

To stick with the boron dopant, but to also mitigate the deleterious effects of boron–oxygen complexes, it is generally helpful to use wafers with the lowest concentration possible for both species [141,142]. To offer a reduced concentration of boron, cell manufacturers could conceivably adopt higher resistivity wafers. But there is a very narrow limit to how far this can be exploited because there is a tradeoff in the form of an increased J_o—and a correspondingly lower V_{oc} and FF—as it becomes increasingly difficult to achieve acceptable ohmic contact resistance between high resistivity silicon and metal electrodes [113,143]. As for how to reduce the concentration of interstitial oxygen, the currently best-known solutions seem to involve alternative approaches to crystal growth—the foremost being either the float zone (*Fz*), or magnetically confined *Czochralski* (M-*Cz*) techniques [70,71,144]. These alternative approaches to crystal growth offer a reduced concentration of interstitial oxygen, either because they do not use a silica-based crucible (*Fz*), or because there is better control of the convection currents within the silicon melt that are exposed to the crucible (M-*Cz*) [30]. Either approach can deliver wafers with lowered oxygen content, and, depending upon the base resistivity, stabilized carrier lifetimes ranging from the hundreds of microseconds to several milliseconds [69,70,135,145]. However, even though the original PERL cell was built upon a *Fz* base wafer, the M-*Cz* option is currently the more commercially relevant because *Fz* ingots are significantly more expensive to produce and have traditionally been limited to a much smaller diameter [146].

5.4. Technology group 2: the interdigitated back contact (IBC) c-Si solar cell (≈ 25% cell efficiency)

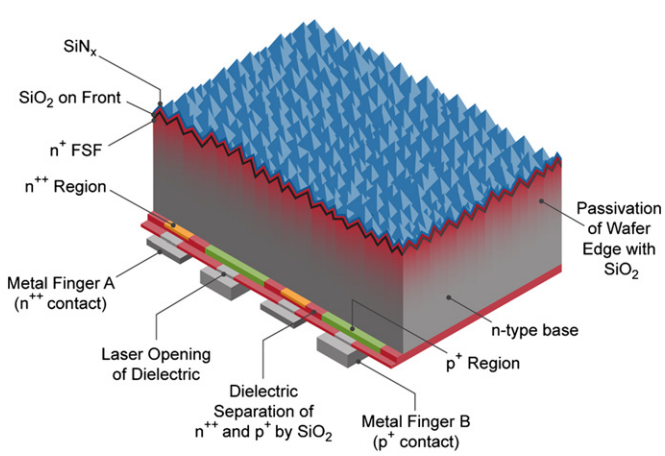

5.4.1. On the benefits and requirements of making a c-Si solar cell with hidden metal electrodes

One of today's more esthetically pleasing PV modules has no obvious metal connections on top of or between the cells. Even to a layperson this design makes sense because there should be increased absorption in the solar cell by eliminating frontside metal grid shading—and indeed the relative gain in J_{sc} to be enjoyed by eliminating these optical shadowing losses is a noticeable 5–10% (depending upon the finger and busbar layout for the front-contacted cell [61]). As another, not-so-obvious benefit of this design, by locating the metal contacts on the back there can also be a much greater emitter-to-metal area coverage. This helps mitigate FF losses, because there can be an overall lower series resistance without also increasing optical shading and reflection losses at the same time. But, in terms of enabling the full efficiency potential of c-Si in commercial-production, there are other advantages to this architecture that are just as hidden as the metal contacts.

In a rear-contacted solar cell, charge carriers generated from the absorption of light near the front surface of the wafer—as most are—must traverse the entire wafer thickness in order to reach the interdigitated array of metal electrode terminals and *p*-/*n*- electrical junctions that are located on the back. With corresponding lifetime requirements in excess of 2 ms for today's typical wafer thickness [65], and with a requirement that the wafers be obtained from industry-relevant manufacturing methods, this currently limits the wafer choice to just one option: those with an *n*-type (phosphorus) base made by the *Cz* approach [54,78,147]. This is because, although the M-*Cz* option can greatly reduce the concentration of problematic boron–oxygen complexes, and the Ga-*Cz* option can offer light-stabilized lifetimes on the order of hundreds of microseconds, largely due to a much lowered sensitivity to residual metallic and carbon impurities it is still the P-doped *Cz* option that most consistently delivers wafers having light- and temperature-stabilized carrier lifetimes on the order of milliseconds [88,141,148].

Because the carrier diffusion lengths for an optimized back-contact solar cell do have to be so high, free charge carriers within these cells will spend more time interacting with all of the wafer's surface areas [80]. This makes consideration of the total SRV even more critical for an IBC cell than for cells made with a lower lifetime base material. Most generally because of the stability of *n*-type wafers under high temperature oxidizing conditions—especially relative to the standard B-*Cz* wafers—this requirement can be easily addressed by utilizing the excellent dielectric surface passivation properties of SiO_2 [79], which affords uniform coverage over the entire wafer front, back, and edge surfaces in one simple thermal processing step. On the front of an IBC cell, after establishing a front surface field via approaches such as light $POCl_3$ diffusion [97,149], a SiN_x:H on SiO_2 stack can provide three advantages: very high quality surface passivation, high transparency to sunlight, and excellent antireflection properties. On the back of the cell, an optically thick layer of Si/SiO_2 on a back metal stack provides very high reflectivity due to the very low refractive index of SiO_2 (n=1.46) [63,64,86,126]. Finally, due to the proximity of the emitter and base contacts within the backside array, in an IBC cell a material that is capable of electrically isolating each n^{++} and p^{+} diffusion should be also be present in order to mitigate current leakage—yet another need that can be met with SiO_2 [88,150].

In consideration of the ability of P-doped *Cz* wafers to use SiO_2 for all of these purposes, and in consideration of the fact that it can be deposited very easily and cheaply in industrially scalable processes, we have incorporated a thermal oxidation step within our process flow used to model the costs of IBC cells (Fig. 8). Most of the concepts and processing steps that we apply are similar in nature to those discussed within the Standard and Technology Group 1 sections. Significantly, however, we have called upon a different procedure (step 5) for the drive-in of additional P dopant atoms in order to make the n^{++} diffusion region.

Because the electrical connections to the base (back surface field) contacts must also be incorporated into the same area that

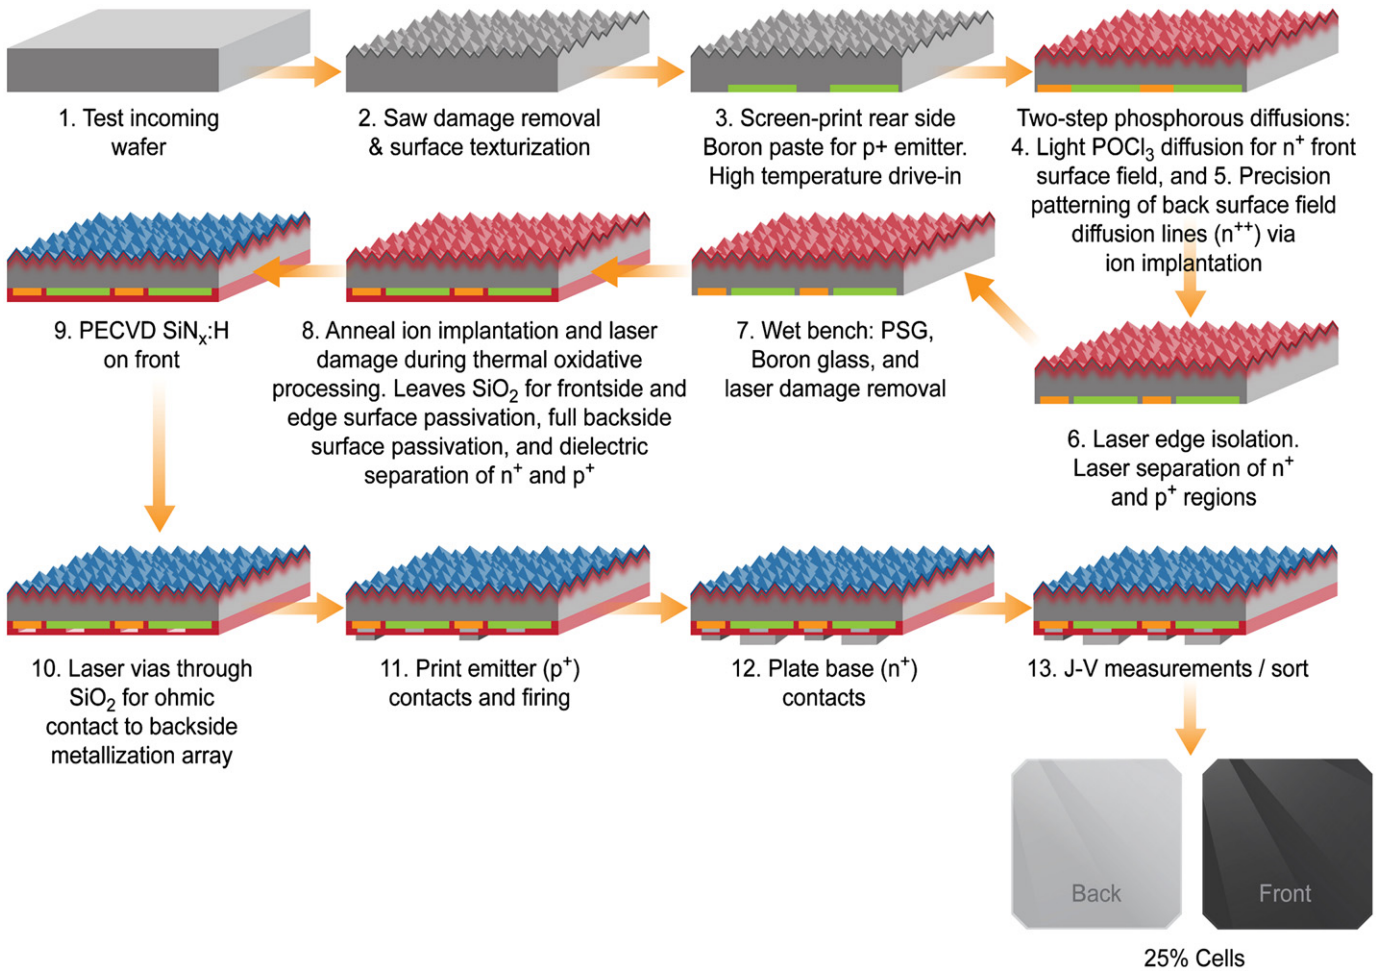

Fig. 8. Modeled process flow for fabrication of the Technology Group 2 (IBC) c-Si solar cells.

would otherwise be fully available for collecting and moving charge carriers to the emitter contacts, the *total* area available for the capture of the all-important minority charge carriers is reduced for an IBC cell relative to a front-contacted cell. For the case of an *n*-type base, this means there is a reduction in the probability for capturing minority carrier holes, before they recombine with majority carrier electrons, because there is less area with which to do the capturing. This is a problem more generally called 'electrical shading' and, if not handled appropriately, it can noticeably compromise the efficiency gains that are enjoyed by an IBC cell by virtue of the fact that it eliminates optical shading losses [76,77].

To address this efficiency loss mechanism, the general goal should be to maximize the total area of the p^+ emitter diffusions while minimizing the total area of the n^{++} base diffusions. Moreover, this must also be achieved within an interdigitated array having base diffusion line widths that are as narrow as possible, and with an optimized spacing [88,151,152]. To achieve the very narrow, very precise, and high aspect ratio base diffusion profiles that may be needed for this optimization, we make an accommodation within our process flow that the n^{++} diffusion in a fully-optimized IBC cell may require something beyond simply exposing wafers to $POCl_3$ and/ or screen-printing a dopant paste. Our cost modeling analysis is based upon a compilation of manufacturing costs related to the dry and in-line process of precision patterned ion implantation [73–75], which has already been commercially developed for the much more technologically complex integrated circuits industry [153], and which is just now being developed to achieve wafer throughputs that are relevant to the PV industry [154].

5.5. Technology group 3: the heterojunction with intrinsic thin layer (HIT) cell (≈ 24% cell efficiency)

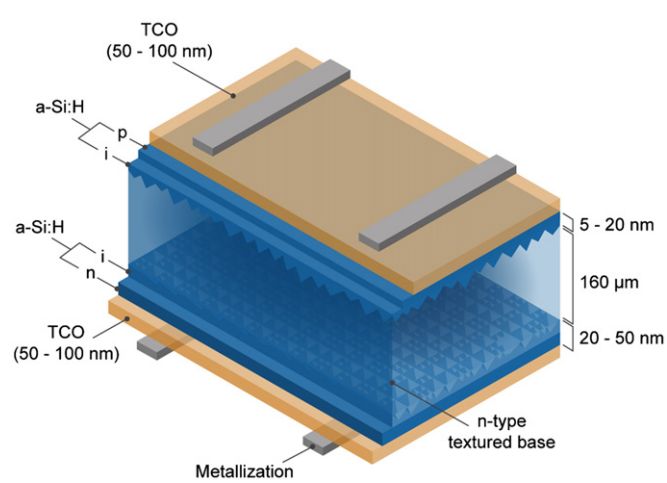

As the final cell architecture considered within this roadmap, we consider cells which utilize very thin a-Si:H layer stacks on

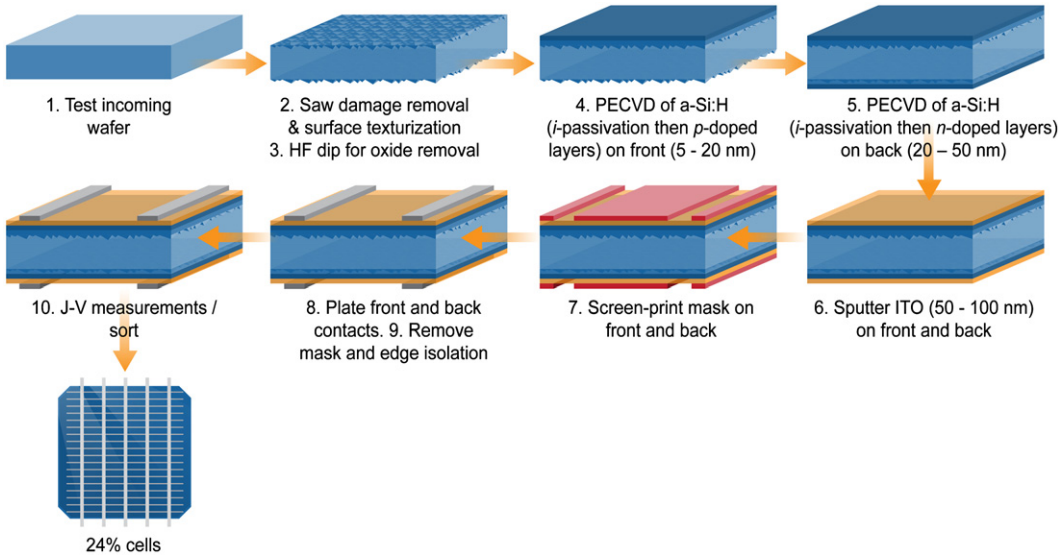

Fig. 9. Model process flow for fabrication of a bifacial Technology Group 3 (HIT) cell.

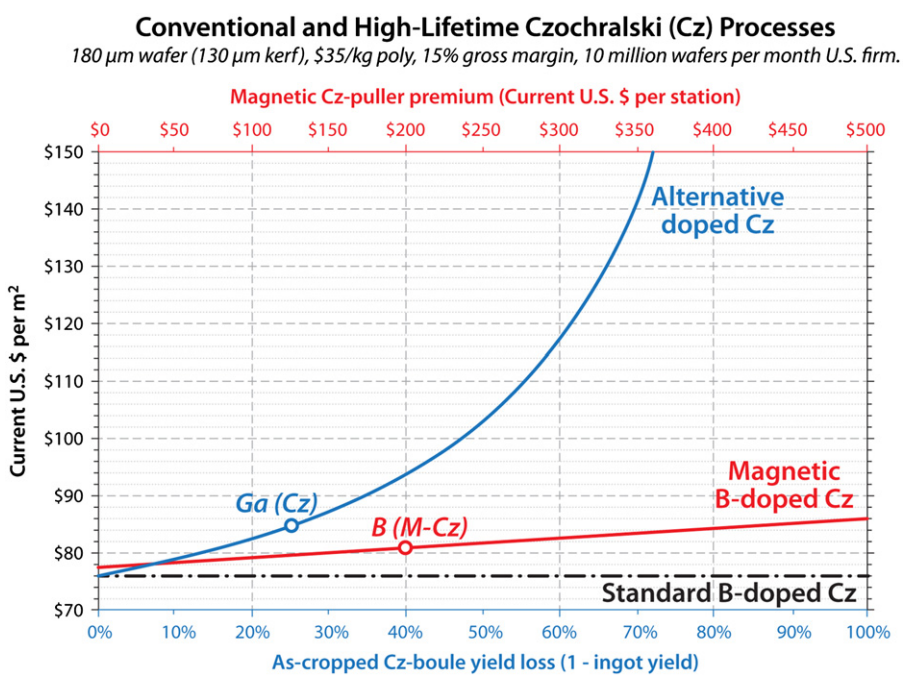

Fig. 10. Sensitivity of minimum sustainable wafer prices to changing ingot wafer yields and initial capital equipment outlay, based upon 1 H 2012 polysilicon prices (blended price of $35/kg) and the wafering cost analysis outlined in Section 4. The horizontal dashed line represents the baseline (180-µm) B-*Cz* wafers shown in Fig. 5.

n-type wafers to provide surface passivation, emitter formation, and a back surface field [81]. Not only have these 'HIT' cells achieved commercial-production efficiencies that are a close second to the IBC cells [155], they can also offer some compelling benefits at the LCOE level as well. First, high V_{oc} HIT cells offer a temperature coefficient that can be slightly lower than the IBC cells, and almost half that of a standard c-Si cell [92]. Second, HIT cells easily offer the possibility to realize bifacial structures, which can lead to greater total harvesting of solar power over a system's lifetime [156]. HIT cells offer another potential benefit in that they can be fabricated using a very simple processing sequence that can be carried out—in its entirety—below 200 °C [112,157].

A typical architecture for HIT cells is shown above, and a representative sequence for fabricating them is shown in Fig. 9. When fabricating these *p*-/*i*- and *i*-/*n*- stacks, the thickness of the intrinsic and doped a-Si:H layers must be carefully optimized with two conflicting requirements. On the one hand, the V_{oc} and FF generally increase as the front layer stack is increased, and then levels off at around 20 nm total thickness. As a tradeoff, because of undesired absorption in the a-Si:H layers, the J_{sc} gradually decreases for all total layer thicknesses greater than around 5 nm [158]. This preference for the minimum allowable a-Si:H layer stack thickness stems from the fact that it is generally more advantageous to allow the sequence of light absorption and charge carrier separation to occur within the crystalline wafer—rather than within the amorphous layer—because the material defect density is significantly higher in a-Si:H [157,159,160]. On the bottom of the wafer, excellent surface passivation is provided by a back surface field that can be formed

by depositing an *n*-doped a-Si:H layer on top of an intrinsic a-Si:H layer, with a combined optimal thickness for these two layers somewhere in the range of 20–50 nm. Without knowing what the actual layer thicknesses need to be on either side of an optimized HIT cell, we have expressed these ranges in Fig. 9 and use the middle values for each within our cost models.

Because the chemical processing sequence that takes place immediately before the a-Si:H deposition is critical to obtaining a high V_{oc}, the exact recipe and processing details of this step in the HIT process flow is one of the more proprietary trade secrets in PV. In laboratory-level research, a wet chemical bath treatment called the 'RCA clean' is typically employed [161]. But this specific cleaning process would likely be too expensive for mass production of solar cells [160]. While not having the exact recipe that is used to fabricate the most efficient HIT devices, we assume that some form of wet bench treatment is nevertheless still necessary immediately prior to the intrinsic a-Si:H layer deposition. We base our cost estimates for this step in the process flow upon the single step of HF oxide removal (step 3), after the description given in [162].

One of the last steps for fabricating a HIT cell is to deposit a transparent conducting oxide (TCO) layer, which is step 6 in our model process flow [160]. The first purpose of this layer is to facilitate better charge carrier collection (because a-Si:H has a very low lateral mobility for free charge-carriers, but a TCO on its surface can greatly assist in shuttling mobile charge carriers to the metal electrodes with low series resistance losses). The second purpose of the TCO is to serve as an AR coating. While fulfilling its dual purposes, the exact thickness and conductivity the TCO layer should be optimized towards the best-possible cell efficiencies. This typically means that consideration of the optical clarity and AR properties leads to very thin layers having a less-than-optimal conductivity (because a 'transparent' electrode generally inhibits the complete passage of light to some degree, and because this parasitic absorption is most significantly influenced by the thickness and the free-carrier concentrations within the layer [163]). While there may be several material options for the TCO, the better-known solutions at present are indium tin oxide (ITO) or zinc oxide doped with either boron or aluminum. While it is worth noting that other deposition techniques such as chemical vapor deposition may also present opportunities for forming TCO materials on HIT cells, our specific process flow calls for ITO deposited by sputtering—mainly to highlight that the calculated costs for using this very thin layer of a precious material are probably lower than one might think.

The metallization of HIT cells can most simply be achieved through the standard process of screen-printing, and indeed this is presumably how it is most frequently done in commercial production. However, the a-Si:H layers cannot be taken to a deposition temperature above 200 °C, and this requirement then excludes the use of the standard screen-printed metal pastes. This requirement for low-temperature metallization can be a significant drawback for HIT cells, as the total amount of the low-temperature paste that is needed is greater because the resistivity is a factor of two to four times higher than the standard pastes. Meanwhile, the price for the low temperature pastes, at around $1700/kg, is also notably higher than the price for the standard pastes, at around $1300/kg (2012 U.S.). These dual factors—metallization price and resistivity—present some severe cost hurdles for HIT cells when the standard process of screen-printing is employed. For this reason, and because we have also called upon the same path for the Technology Group 1 and IBC cells, towards the end of our HIT process flow we call upon the lower cost metallization process of electroplating a Cu-alloy on top of an electroless seed of Ni. To prevent uniform plating of metal over the *entire* surface of the TCO, we have also included an additional low-cost patterning step in the process [164]. For our purposes this is assumed to be screen-printing of a resist mask, followed by its removal in an industry standard solution.

The next reasonable question then becomes the cost for the electroplating sequence relative to the cost of screen-printing copious amounts of expensive silver. In the next section, we address this question and other questions that might naturally arise for the different cell designs that have been presented.

5.6. Cells cost analysis: the overall results and some specific points of note

5.6.1. On the potential price premium for higher lifetime wafers

To begin a discussion on costs, because the two are so intimately related it is first necessary to provide some necessary disclaimers on the assumed efficiencies. A significant assumption underlying the efficiency projections shown in Table 3 is that the higher efficiency cells must be built upon a foundation of higher quality starting material than the standard B-*Cz* wafers. In total, a full 1–2% improvement in the absolute efficiency, above the standard cell, may be possible solely due to the utilization of higher lifetime wafers [88,135].

The conventional thinking is that, in order to provide higher lifetime wafers to cell manufacturers, for one of two reasons there should be an attached price premium—either because wafer manufacturers must charge more because they will have suffered higher yield losses in the alternative dopants case, or because of additional capital equipment and energy costs if they were to attach magnetic systems to *Cz* pullers.

Allow first for some explanation of what is meant by 'higher yield losses in the alternative dopants case'.

The traditional boron dopant, with a relatively high segregation coefficient of 0.8 [42], has an inherent propensity to distribute very well during the process of pulling a *Cz* boule. This then directly leads to variations in base resistivity that are typically very low, at less than 0.5 Ωcm, for all wafers taken along the entire length of the boule [135,144]. In the absence of additional sorting steps to more carefully bin wafers by their base resistivities, by virtue of their being taken from different parts of a boule, this low variation enables a corresponding consistency in current and voltage for essentially all cells made from all as-received B-*Cz* wafers. The desire to have this consistency is certainly clear (it is, to begin, crucial for reducing the cell-to-module derate), and it is in this respect that the yield loss for wafers made with the B dopant is nominally zero.

But not everything is ideal for the B-*Cz* wafers because they are, to one degree or another, vulnerable to problems associated with metastable boron–oxygen complexes. One demonstrated solution to this problem is to use alternative dopants—such as Ga for making *p*-type wafers or P making for *n*-type wafers. With lower segregation coefficients, however, these alternative dopants do not distribute as uniformly as boron does—leading directly to greater axial variations in base resistivity, and, potentially, corresponding yield losses because not every wafer is usable to cell manufacturers. Quantifying these yield losses would require knowing the average number of wafers per ingot that are unsellable, based upon how the average resistivity varies for wafers taken from different sections of all produced ingots, and exactly what range of base resistivities cell manufacturers are interested in buying. For the standard cell made with *p*-type wafers, the allowable range of base wafer resistivity is generally anywhere between 0.5 and 3 Ω cm [142]. While such a range is easily achieved with the boron dopant, this desired resistivity range would correspond to a wafer yield loss of around 25% for a Ga-doped *Cz* ingot produced by single charge loading and without semi-continuous feeding [144,165].

To partially offset this material yield loss in wafering, the head, tail, and chord scrap of an ingot made with a low segregation coefficient dopant can still be recycled for later rounds of *Cz* ingot pulling. While incorporating this partial material recovery into the wafering cost model, in Fig. 10 we display the sensitivity of minimum sustainable wafer prices to the material yield loss in the wafers-only section of the ingot. For a 25% wafer loss, by our analysis a maximum increase in wafer price of around 10% would be necessary.

Since Phosphorus has a much higher segregation coefficient than Gallium (0.35 versus 0.008 [135]), a less than 25% yield loss would be expected for the P-doped *Cz* wafers. As an additional consideration for selling such *n*-type wafers, the allowed base resistivity can actually be unique for each c-Si cell manufacturing process—with a much broader range of 2–10 Ω cm allowed for cells made on an *n*-type base [78,149,166]. In that case, such an allowable range would translate to no yield loss, and therefore no necessary price premium, for the doping of *Cz* ingots with Phosphorous. Of course, this explanation does not exactly match the historical pricing trends where there was a price premium for *n*-type wafers. But it is a consensus echoed by every interviewed industry collaborator: any price premium that may have existed for the *n*-type *Cz* wafers was simply a result of their being a product for a 'niche market', and not because there is some upstream loss in wafer yield.

For the purpose of characterizing the costs associated with the magnetically confined *Cz* process, the lower curve in Fig. 10 corresponds to the relationship between minimum sustainable wafer prices and any additional capital equipment and energy costs that may exist for the *Cz* pulling of ingots. The baseline wafer price, equal to \$76/m^2, was derived in Fig. 5. The first calculated small revision upwards in cost—the higher y-intercept of around \$2/m^2—is due to the additional energy requirements of a typical M-*Cz* system. This price correction is based upon an aggressive assumption that 1980 kWh of additional electricity would be required for each 2.158 long, 205 mm diameter boule that is produced (and detailed in Section 4).

With respect to the additional capital equipment costs for the M-*Cz* attachment, the necessary wafer price is estimated to have a linear dependence. From a survey on crystal growth equipment, we gathered that the applicable M-*Cz* pullers came with an additional median cost of around \$200,000 (current U. S.), although it could be twice that much for the more expensive option that employed superconducting magnets [146,167]. By our sensitivity analysis the additional capital equipment costs over that range would translate to an additional price premium of \$3–\$6/m^2 of wafers produced. At 20% cell efficiency, this would translate to a cost penalty of just \$0.01–\$0.03/W.

Ultimately, all signs are that any additional costs that might be associated with fabricating higher lifetime wafers are probably quite small. For making *p*-type wafers, it would also appear to be a difficult decision between using an alternative dopant such as Gallium—with its potentially compromised wafer yield—versus making a bigger investment in the capital equipment and energy inputs for M-*Cz*. (Although, for any yield loss expected to be greater than around 20%, there may be an advantage in paying that little extra for the M-*Cz* upgrade—instead of being left to wonder just what to do with all those boxes filled with unsellable wafers.)

5.6.2. Some statements on the overall cells cost analysis

In Fig. 11, the results are shown for the projected manufacturing costs of the Technology Groups 1–3 cells, with the cost-of-ownership models for each step of the process flows shown in Figs. 7–9. (Please see the supplementary information for a breakdown of costs for each step). In the long-term case it is assumed that the *n*-type wafer price is \$22/ m^2, corresponding to the
80-μm kerfless wafers shown in Fig. 5. That is, it is assumed that there will be no long-term price premium for those wafers due to potential yield losses in wafering. In order to attach the potential price premium for higher lifetime *p*-type wafers, the wafering cost model was rerun at the 80-μm thickness. The curves for those thinner wafer curves were similar to Fig. 10 and showed a marginal penalty of just \$1–\$2/m^2 for the M-*Cz* option. This less than \$0.01/W penalty is included within the projected costs for the Technology Group 1 cells shown in Fig. 11. For all cell types, the results shown assume that the yields in cell manufacturing are constant—even for the case of thinner wafers.

The IBC cells appear to have the highest total expected depreciation expense; however, the difference becomes significantly reduced for similarly sized wafers. This is because the modeled process flows for the Technology Groups 1 and 2 are actually very similar, and because the effective throughput of the cell manufacturing equipment is estimated to be lower for smaller area wafers (if equipment is sold based upon throughputs measured in the number of wafers per hour, then for a given production volume more pieces of equipment must be purchased in order to handle more smaller area wafers). This is why, for the sake of comparison, we include a case that might be made for the Technology Group 1 cells to be fabricated on 160 μm wafers.

In further regards to the depreciation expense for the IBC cells, we also calculate that the capital costs for the ion implantation tools may not be as significant as one might think. With a wafer throughput of roughly 1100 wafers per hour for a roughly \$1.5 million machine, by our calculations ion implantation contributes less than \$0.02/W to the total manufacturing cost for our model IBC cells.

In light of the compelling efficiency for the HIT cells, and by virtue of the fact that they have a very simple manufacturing process, it comes as no surprise that they are also calculated to have a very competitive manufacturing cost. But that calculated result is contingent upon the replacement of the screen-printed, low-temperature Ag pastes with an electroplated Cu alloy.

Because the processing temperatures for HIT cells must generally be kept below 250 °C, the cocktail of binders within the low-temperature metallization pastes differ from the binders that can be used in the standard metallization pastes (which can be made from materials that are more easily removed at the typical 800–900 °C temperatures endured during a standard co-firing step [160,163]). This currently leaves the low-temperature metallization with a compromised resistivity (typically between 15–25 mΩ per square for 25 μm of planar paste thickness) relative to the standard Ag pastes (typically between 4 and 5 mΩ per square for 10 μm of planar paste thickness) [168]. To overcome the higher resistance within the low-temperature paste, but to still keep the I^2R losses constant, there are several approaches that can be employed: use gridlines with a greater cross-sectional area; use more tightly spaced gridlines (although this comes with a trade-off in that more busbars are needed, and in optical shading), and use smaller diameter wafers.

By our calculations, to achieve a similar resistivity it is still necessary to deposit around four times as much low-temperature paste per unit area of manufactured cells—even when using 155 cm^2 wafers instead of 237 cm^2 wafers, and even when employing an optimized metallization geometry. Considering the \$1700/kg price for the low-temperature paste versus \$1300/kg for the standard paste, and with a total baseline material requirement of around 200 mg of Ag paste for the frontside grid and rear busbars for every 17% efficient 237 cm^2 cell made in the standard process [169,170], a quadrupling in the area-based paste requirement would

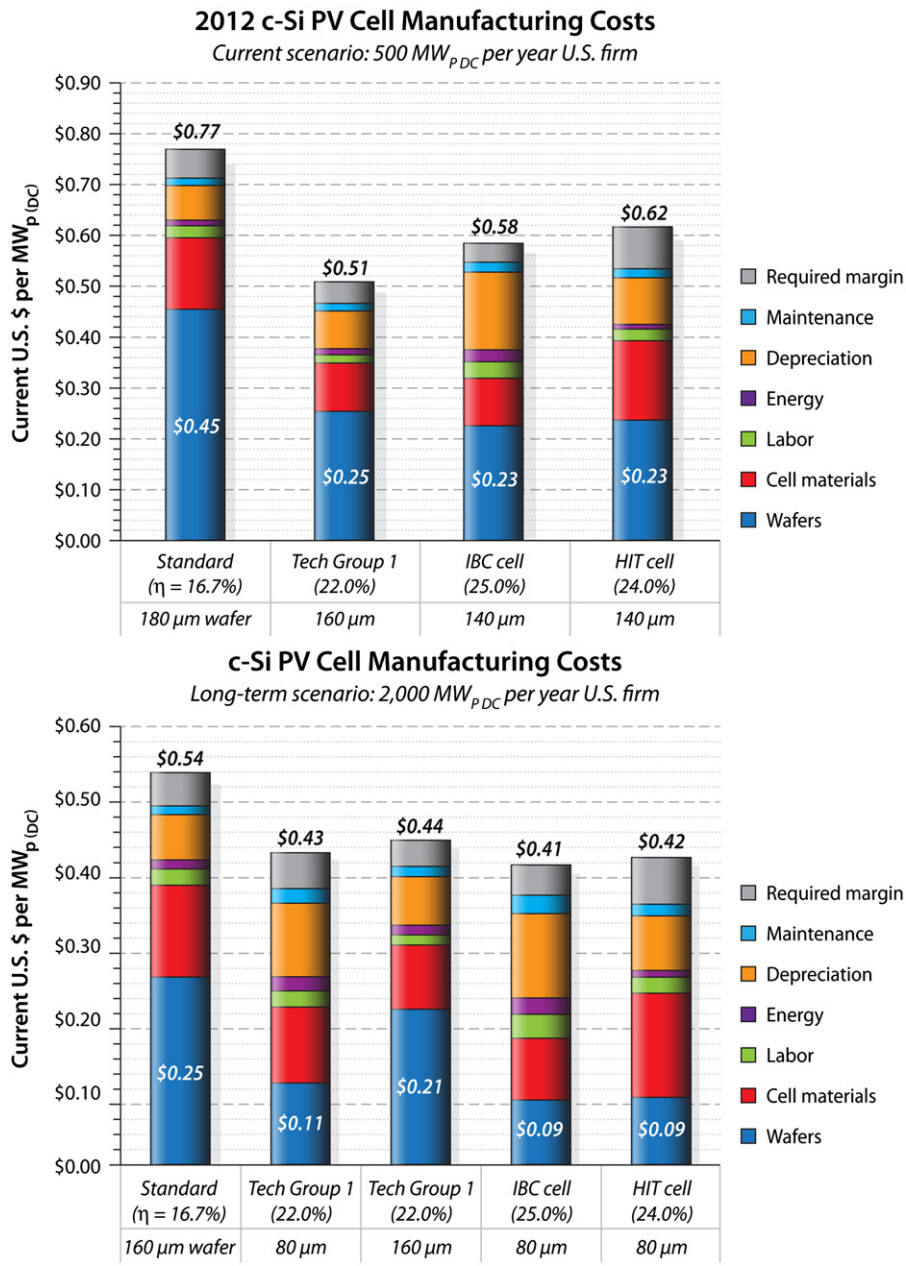

Fig. 11. Cost model results for cells for the Standard and Technology Groups 1–3 cells. Top: cell costs derived with typical 2012 wafer thicknesses (shown at the bottom of each bar), and the wafer prices depicted in Fig. 10. Bottom: estimated cell-processing costs with either the 80-μm or the 160-μm wafers, where all cells are made from $23/kg polysilicon and the future-case wafer price premiums mentioned in the text. The WACC used for assigning the required margin is 8.6% in the 2012 scenario and 6.2% in the long-term scenario.

then work out to a metallization cost disadvantage of around $0.17/W for the 24% efficient 155 cm^2 HIT cells.

This certainly helps to explain why the HIT cells in particular could benefit from alternative metallization schemes. With the requirement that any alternative metallization scheme must still be a low-temperature process, it is perhaps logical that it would be most beneficial for HIT cells to call upon the process of electroplating. This process was also used as a metallization procedure in the Technology Group 1 and IBC cell models; but, in the case of the HIT cells there would have to be an additional masking step in order to carry out the electroplating within a precise pattern (otherwise, without a mask, there would be metal deposited over the *entire* surface of the TCO). The relevant cost-of-ownership model that we were provided for the screen-printing of a masked resist would be expected to add around $0.01/W for both the front and rear sides. Presumably there should also be a mask removal step, but—if it can be done with an industry standard resist stripping solution—this is also a conceivably very cheap step (so cheap, in fact, that even in tandem with chemical edge isolation the two-step wet bench process would only cost around $0.02/W total).

To conclude the discussion on cells, several of the cell processing flows referenced the PECVD deposition of SiN_x:H, Si:C, or a-Si:H layers. The sputtering of ITO was also mentioned. For all of these layers, we find that the depreciation expense is relatively consistent over the range of projected thicknesses. This is because the total cycle times depend more upon the longer steps of wafer handling and pump-down than the total times required for depositing even the thickest layers. We also find that the sputtered ITO layer in the HIT cell is more affordable than one might think. Primarily because

it is a very thin layer, our estimated net material costs for sputtering 75 nm of ITO from a rotary target are less than \$0.02/W for each side of the bifacial cell. Knowing this, and knowing that the electroplating of Cu on a HIT cell has recently been demonstrated [164], from a costs standpoint it would be a pretty tight contest between all three of these modeled c-Si solar cell architectures.

6. Modules

In the final step of the c-Si supply chain, completed cells are incorporated into modules by first electrically connecting cells together into strings with conductive solder and tabbing ribbons. The ends of those strings are then soldered onto bussing ribbons. To protect this assembly from the elements, it is encased within a top-bottom stack of encapsulant films—typically ethyl vinyl acetate, or 'EVA'—that have been melted ($T_{melting} \approx 145–160\ ^{\circ}C$) and vacuum-laminated onto the array. During this encapsulation step, the assembly is also bonded to a sheet of front glass and to a backside film or glass with a tape that is dispensed around the perimeter of the module. An aluminum frame is also oftentimes fit around the perimeter of the module—with the benefits that it can be used to protect the module edges, to provide a connection point for electrical grounding, to support snow and wind loads, and to make the module installation an overall easier process. (The frame is, however, a relatively expensive component, and there is still an open debate within the industry for how to realize those same benefits at a reduced cost.)

The array of bussing ribbons connected to the ends of each series of strings is then crimped towards a through hole in the module backsheet film. The bussing ribbons are connected to bypass diodes, which are housed inside an electrical junction box, or 'J-box'. The purpose of these bypass diodes is to prevent excessive reverse current flow and power consumption in cells that may be receiving different amounts of sunlight, such as when the module might be partially covered with snow, dust, or leaves, or by some other obstruction [171]. As a final step in module assembly, the 'J-box' is set in place with adhesive sealant on the bottom of the backside film.

The materials and equipment for this final step of the supply chain have become fairly standardized over the years. In Table 4, we detail the typical costs for these materials in early 2012, roughly averaged over numerous conversations with c-Si module manufacturers and their primary material suppliers. In consideration of the optical properties that may be needed within the front glass in order to achieve the target module efficiencies (i.e., in an effort to reduce the amount of parasitic light absorption from certain contaminants such as Fe), and to ensure the mechanical robustness of modules, we have assumed that the highest quality front glass might be needed. There are certainly cheaper types of glass that are available, however. (Please see the supplementary information for other possible front glass cost assumptions).

Within the industry, an intensive effort is underway to identify lower cost module materials and processes. But the adoption of these new approaches is tempered by a very clear need to maintain product bankability. This makes it unlikely that these materials will be significantly changed for at least the foreseeable future. Over the long-term, however, it is possible that the movement to thin or ultrathin wafers may necessitate that the final module materials be modified or even incorporated into wafer handling and cell processing, as many of the ideas that have surfaced for reducing wafer thickness frequently hinge upon the need to use the final module materials as a mechanical support and/or electrical conduit for the more delicate wafers. For example, the front glass/encapsulation combination may need to serve as an adhesive support for wafer bonding and cleavage from an epitaxial substrate; or a conductive film, paste, or epoxy may be needed to electrically connect very thin cells, should the stresses of conventional tabbing and stringing prove to be devastating [172–175].

Without knowing the exact characteristics and purposes of these next-generation module materials, it is correspondingly, difficult to speculate on what their associated costs might be. Thus, as things stand, there is little choice but to assume the same balance of module materials costs shown in Table 4 within the long-term module price projections shown in Fig. 12. With the balance of module materials constant across all technologies, it is the difference in cell efficiencies that explains the final—very subtle—differences in costs in dollars-per-watt terms. Within the figure, the minimum required margins for meeting a 6.2% WACC are also included for each cell type, and this is why the costs shown in Table 4 are slightly lower than the final minimum sustainable module price numbers seen within Fig. 12.

Table 4
Balance-of-materials costs for c-Si module manufacturing, aggregated from several suppliers for each. The module dimensions are assumed to be 1.65 m × 1.20 m for the large-area (237 cm^2) cells used for the 180- or 160-μm thick Standard and Tech Group 1 designs (which corresponds to 72 cells per module), and 2.26 m^2 for the small-area (155 cm^2) cells used for the 140- or 80-μm thick Tech Groups 1–3 designs (which corresponds to 128 cells per module).

Balance of material costs for modules	
Material	**First half 2012 costs**
Long-term cell prices (See Fig. 11)	\$0.77–\$0.41/W_p
Stringing and tabbing ribbons, metal solder and busing ribbons	\$2.50/module
J-box containing the bypass diodes	\$5.00/module
J-box sealant, bonding tape, printed module sticker label and bus bar covers	\$1.50/module
Aluminum frame	\$20/module
EVA (2 sheets needed)	\$3.50/m^2 for each sheet
Backsheet film (Polyvinyl fluoride, and/ or UV- and chemically-stabilized polyethylene terephthalate)	\$8/m^2
Premium front glass: 3.2 mm, low [Fe], tempered, with AR coating	\$16/m^2
Estimated module materials costs	**\$1.08–\$0.61/W_p**

7. Conclusions

The analysis described in this paper is limited in scope to the *Cz* approach for making crystalline silicon ingots and the wire sawing of wafers. As a major step for reducing costs within this infrastructure, the complete elimination of kerf loss could prove to be a very fruitful endeavor, yielding as much as \$15/m^2 in savings. One possibility for wafer manufacturers to realize these savings may be through the recycling of kerf after switching from standard-wire sawing to diamond-wire sawing. Within this scope of c-Si manufacturing, there is also a generally held boundary condition that the wafer thickness must be 80-μm or greater [41]. If manufacturers at all steps of the supply chain can successfully operate at this wafer thickness, at minimum sustainable polysilicon prices we calculate a cost benefit of around \$9/m^2 of wafers produced, relative to the materials cost for producing today's standard 180-μm wafers. But the calculus of wafer thickness must also consider that the mechanical yields for all steps in the supply chain can vary as a function of thickness, that there might have to be modifications to the capital equipment at each step, and that there can be problems with wafer bowing. As there have been historical disadvantages associated with each of these considerations as the wafer thickness is reduced, and knowing that the total potential savings for the 180- versus 80-μm wafers works out to just \$0.04–\$0.05/W, in the

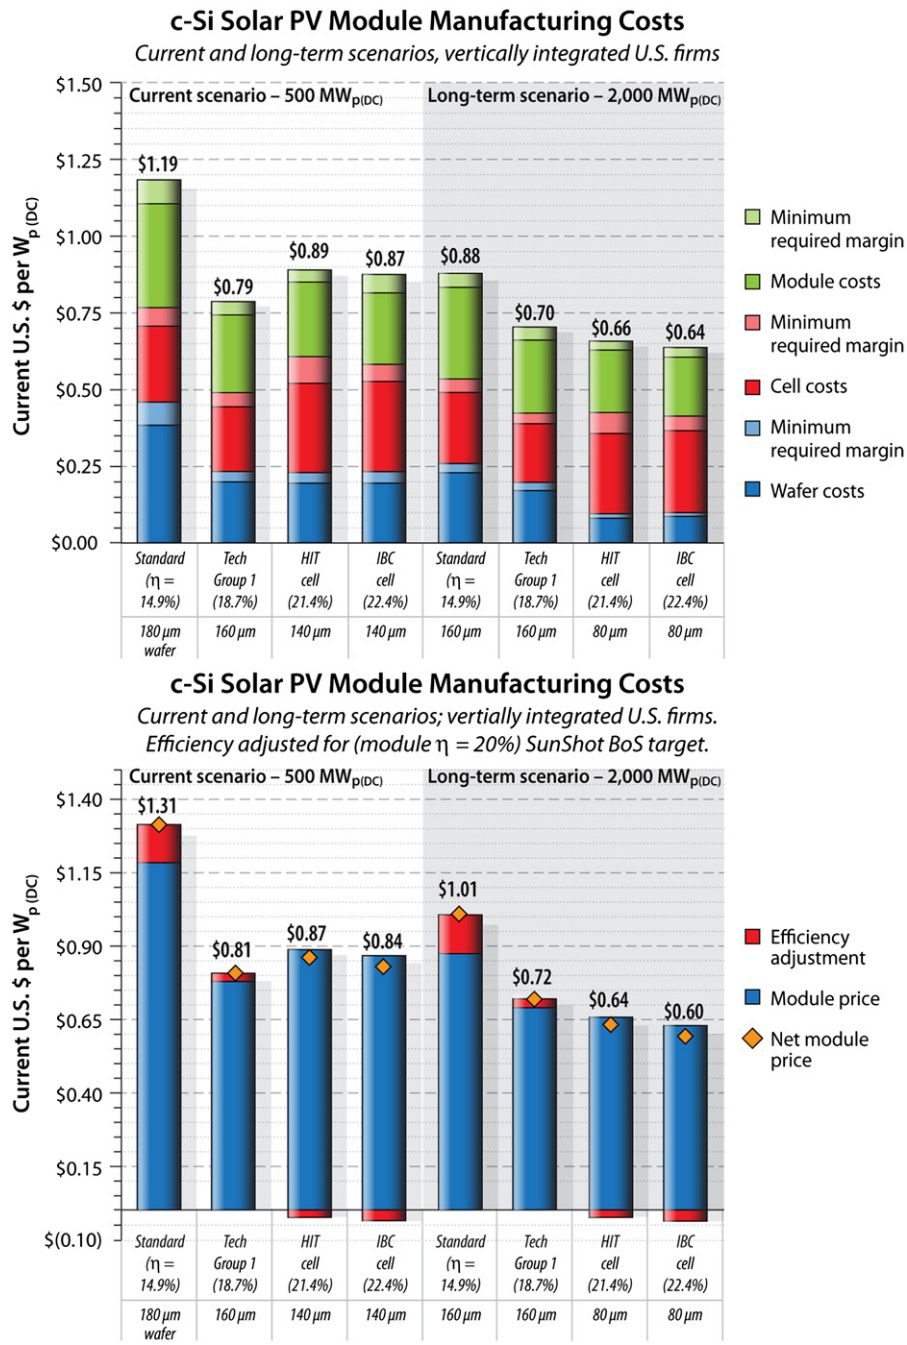

Fig. 12. Top: Cost model results for completed modules: a compilation of estimated costs for manufacturing standard modules and advanced modules within the full c-Si supply chain, assuming all products are transferred at minimum sustainable prices. The numbers underneath each cell type indicate the assumed module efficiency and wafer thickness for each. The long-term scenario reflects the projected costs and prices for modules made with cells on 80-μm wire-sawn kerfless wafers at minimum sustainable polysilicon prices (see Figs. 5 and 11). Bottom: efficiency-adjusted module prices for the different cell types, in consideration of balance-of-systems savings (HIT and IBC) or costs (Standard and Technology Group 1). For a rationale of these efficiency adjustments, please see Ref. [176]. The BOS efficiency adjustments to the module prices are normalized against the 20% module efficiency targeted within the U. S. Department of Energy's SunShot Initiative [177].

end it is still unclear just how rapidly such drastically thinner wafers will be adopted.

We have discussed some of the available opportunities for moving standard c-Si solar cells toward higher sunlight power conversion efficiencies. The advanced cell architectures needed to achieve these higher efficiencies would likely require a greater initial capital equipment expenditure and higher materials costs on a piecemeal basis; but by our calculations the resulting efficiency improvements could very well translate to lower total module manufacturing costs on a dollars-per-watts basis. With three advanced cell architectures in hand, and with best-case wafer prices, we project that c-Si modules made from wire-sawn wafers within the United States could conceivably move from 14.5–22% efficiency and $1.10–$1.45/W minimum sustainable prices at the beginning of 2012 to 19–23% efficiency and $0.60–$0.70/W prices over the long-term. This estimate supports the rigorous derivation of future c-Si module prices from experience-based learning curves—where the 'long-term' price potential was estimated by Nemet to be around $0.65/W [6]. But, while we deliberately shy away from attaching specific dates to our guess of when the 'long-term' module price might be equal to the actually sustainable module price, with a globally strong demand for more PV deployment, and with all of the research and development occurring

in c-Si, it may very well prove to be long before the 'optimistic' date of 2027 derived from his curve labeled the 'aggressive scenario'.

Acknowledgments

We would like to acknowledge the following individuals for helpful comments, corrections, and contributions during the preparation of this manuscript:Allen Barnett (The Universities of Delaware and New South Wales), John Benner (Stanford University), Tonio Buonassisi (The Massachusetts Institute of Technology), Denis De Ceuster (TetraSun), Charles Gay (Applied Materials), Martin Green (The University of New South Wales), Steven Hegedus (The University of Delaware), Stefan Köstner (Max Planck Institute), Sarah Kurtz (NREL), Minh Le (The U.S. Department of Energy), Margaret Mann (NREL), Robin Newmark (NREL), John Lushetsky (The U. S. Department of Energy), Douglas Powell (The Massachusetts Institute of Technology), Hans J Queisser (Max Planck Institute), Ramamoorthy Ramesh (The U. S Department of Energy and The University of California, Berkeley), Doug Rose (SunPower), Oliver Schultz-Wittmann (TetraSun), Ron Sinton (Sinton Instruments), and Richard Swanson (SunPower).

We would also like to thank our many industry collaborators who provided the cost data that made this manuscript possible, as well as those who served for SOLMAT as anonymous reviewers. SolMat Editor: Greg Smestad, Graphics Artis): Alfred Hicks, Technical Publications Editor (NREL): Kendra Palmer.t (NREL).

Appendix A. Supporting information

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.solmat.2013.01.030.

References

[1] P. Frankl, S. Nowak, M. Gutschner, S. Gnos, T. Rinke, International Energy Agency Technology Roadmap: Solar Photovoltaic Energy, Available Online at: 〈http://www.iea.org/papers/2010/pv_roadmap.pdf〉, 2010.
[2] P. Mints, Photovoltaic Manufacturer Shipments, Capacity and Competitive Analysis 2010/2011, Navigant Consulting Photovoltaic Service Program: Report NPS-Supply 6, Palo Alto, CA, 2011.
[3] Sunpower, SunPower E19/320 Solar Panel Datasheet. Available online at: 〈http://us.sunpowercorp.com〉, 2011.
[4] D.D. Smith, P.J. Cousins, A. Masad, S. Westerberg, M. Johnson, X. Tu, T. Dennis, G. Harley, G. Solomon, S. Rim, M. Shepherd, S. Herrington, M. Defensor, A. Leygo, P. Tomada, J. Wu, Sunpower's generation III high-efficiency lower-cost technology: transition to full-scale manufacturing, in: Proceedings of the Thirty eighth IEEE PVSC, 2012.
[5] Trina Solar Linear Performance Warranty (3.5% in the first year, thereafter 0.68% per year, ending with 80.18% in the 25th year after the Warranty Start Date). Available online at: 〈http://www.trinasolar.com/eu/products/product-resources/product-warranty〉.
[6] G.F. Nemet, Beyond the learning curve: factors influencing cost reductions in photovoltaics, Energy Policy 34 (2006) 3218–3232.
[7] D.M. Powell, M.T. Winkler, H.J. Choi, C.B. Simmons, D.B. Needleman, T. Buonassisi, Crystalline silicon photovoltaics: a cost analysis framework for determining technology pathways to reach baseload electricity costs, Energy and Environmental Science 5 (2012) 5874–5883.
[8] J. Chase, Race to the bottom for polysilicon costs, Bloomberg New Energy Finance, Q1 2012 Solar Research Note, 2012 1–8.
[9] Polysilicon prices hit record low in 2011; will head even lower enabling $0.70/W PV in 2012, in: Greentech Solar, February 10, 2012. Available online at: 〈http://www.greentechmedia.com/articles/read/polysilicon-prices-hit-record-lows-in-2011-will-head-even-lower-enabling-0/〉.
[10] 〈http://pvinsights.com/〉 last accessed July 15, 2012.
[11] M. Bazilian, I. Onyeji, M. Liebreich, I. MacGill, J. Chase, J. Shah, D. Gielen, D. Arent, D. Landfear, S. Zhengrong, Re-considering the economics of photovoltaic power, Bloomberg New Energy Finance White Paper. Available online at: 〈https://http://www.bnef.com/WhitePapers/download/82〉, 2012.
[12] C. Breyer, A. Gerlach, Global overview on grid-parity event dynamics, in: Proceedings of the Twenty fifth EUPVSEC, Valencia, Spain, 2010, pp. 5283–5304.
[13] K. Aanesen, S. Heck, D. Pinner, Solar power: darkest before dawn, McKinsey & Company industry report, Available online at: 〈http://www.mckinsey.com/client_service/sustainability/latest_thinking/solar_powers_next_shining〉, 2012.
[14] M. Woodhouse, A. Goodrich, T. James, D. Feldman, T. Markel, R. Margolis, Economic analysis of photovoltaics versus traditional energy sources: where are we now and where might we be in the near future, in: Proceedings of the Thirty seventh IEEE PVSC., NREL Report Numbers PR-6A20-52311 and CP-6A20-50714. Available online at: 〈https://http://www.nrel.gov/docs/fy11osti/52311.pdf〉 and 〈https://http://www.nrel.gov/docs/fy11osti/50714.pdf〉, 2011.
[15] J.V. Busch, Technical Cost Modeling of Plastics Fabrication Processes, Ph.D. Thesis, Massachusetts Institute of Technology, 1987.
[16] International Labor Comparisons (U. S. Bureau of Labor Statistics) &International Industrial Electricity Price Index (U. S. Energy Information Administration), Available Online at: 〈http://www.bls.gov/fls/〉 and 〈http://www.eia.gov/emeu/international/elecprii.html〉.
[17] S.A. Ross, R. Westerfield, B.D. Jordan, Fundamentals of Corporate Finance, McGraw-Hill, Irwin, New York, NY, 2009.
[18] T. Ogier, J. Rugman, L. Spicer, The Real Cost of Capital, Prentice Hall Financial Times, London, 2004, pp. 29–60.
[19] S. Markle, D. Shackelford, Do multinationals of domestic firms face higher effective tax rates?, National Bureau of Economic Research, Cambridge, MA, 2009.
[20] A.C. Goodrich, D.M. Powell, T. James, T. Buonassisi, M. Woodhouse, Benchmarking the competitive manufacturing costs of wafer-based crystalline silicon PV components in the U.S. and China, (In preparation).
[21] A. Goodrich, T. James, M. Woodhouse, Solar PV, Manufacturing cost analysis: U. S. competitiveness in a global industry, Presentation to the Stanford University Precourt Institute for Energy, NREL Presentation report PR-6A20-53938. Available online at: 〈http://www.nrel.gov/docs/fy12osti/53938.pdf〉, October 10, 2011.
[22] 〈http://finance.yahoo.com/〉.
[23] Cost of Capital by Sector 〈http://w4.stern.nyu.edu/~adamodar/New_Home_Page/datafile/wacc.htm〉.
[24] J. Gjerdrum, N. Shah, L.G. Papageorgiou, Transfer prices for multienterprise supply chain optimization, Industrial and Engineering Chemistry Research 40 (2001) 1650–1660.
[25] B. Ceccaroli, O. Lohne, Solar Grade Silicon Feedstock in Handbook of Photovoltaic Science and Engineering, second ed., John Wiley & Sons, Ltd, West Sussex, United Kingdom, 2011.
[26] B. Hazeltine, C. Fero, W. Qin, Advancements in the commercial production of polysilicon, Photovoltaics World 3 (2010) 16–22.
[27] A. Tour, M. Glachant, Y. Meniere, Innovation and international technology transfer: the case of the Chinese photovoltaic industry, Energy Policy 39 (2011) 761–770.
[28] G. Zhu, H. Shu, T.H. Chau, GCL-Poly Energy Holdings Limited: 2011 full year results, Available online at: 〈http://www.gcl-polycom.hk/attachment/2012031518371817_en.pdf〉, 2012 1–25.
[29] W.O. Filtvedt, M. Javidi, A. Holt, M.C. Melaaen, E. Marstein, H. Tathgar, P.A. Ramachandran, Development of fluidized bed reactors for silicon production, Solar Energy Materials and Solar Cells 94 (2010) 1980–1995.
[30] S. Kasap, P. Capper, Single-crystal silicon: growth and properties, Springer Handbook of Electronic and Photonic Materials, Springer Science and Business Media Group, Leipzig, Germany, 2006 (Chapter 13).
[31] R.H. Hopkins, R.G. Seidensticker, J.R. Davis, P. Rai-Choudhury, P.D. Blais, J.R. McCormick, Crystal growth considerations in the use of solar grade silicon, Journal of Crystal Growth 42 (1977) 493–498.
[32] S. Middleman, A.K. Hochberg, Process Engineering Analysis in Semiconductor Device Fabrication, McGraw Hill, New York, 1993.
[33] Applied Materials whitepaper: wafer wire sawing economics and total cost of ownership optimization. Available online at: 〈http://www.appliedmaterials.com/〉Wire-Sawing-Economics-and-COO-, 2011.
[34] S.K. Chunduri, Diamond wire—the die has been cast, Photon International 5 (2011) 258–263.
[35] M. Schumann, T.O. Perez, S. Riepe, The solar cell wafering process, Photovoltaics International 5 (2009) 53–59.
[36] S.K. Chunduri, Diamonds—the cutting edge? Market survey on sawing and grinding equipment, Photon International 5 (2011) 222–256.
[37] B. Prior, C. Campbell, Polysilicon 2012–2016: supply, demand and implications for the global PV industry, GreenTech Media Research 1 (2012) 1–72.
[38] D. Kray, H. Kampwerth, E. Schneiderlochner, Comprehensive experimental study on the performance of very thin laser-fired high efficiency solar cells, in: Proceedings of the 19th EUPVSEC, Paris, 2004, pp. 608–611.
[39] A. Wang, J. Zhao, S.R. Wenham, M.A. Green, 21.5% efficient thin silicon solar cell, Progress in Photovoltaics: Research and Applications 4 (1996) 55–58.
[40] D. Kray, K.R. McIntosh, Analysis of ultrathin high-efficiency silicon solar cells, Physica Status Solidi A 206 (2009) 1647–1654.
[41] E.V. Kerschaver, K. Baert, J. Poortmans, Challenges in producing photovotaic modules on thin wafersfourth ed., Photovoltaics International 64–67.
[42] S. Middleman, A.K. Hochberg, Process Engineering Analysis in Semiconductor Device Fabrication: Silicon Production, McGraw-Hill Inc, New York, 1993, pp. 230–259.
[43] F. Dross, K. Baert, J. Poortmans, Crystalline silicon thin foils: where crystalline quality meets thin-film processingfifteenth ed., Photovoltaics International, pp. 107–113.

[44] D. Neuhaus, A. Munzer, Industrial silicon wafer solar cells, Advances in Optoelectronics (2007) 24521.
[45] R.A. Sinton, A. Cuevas, Contactless determination of current-voltage characteristics and minority carrier lifetimes in semiconductors from quasi-steady-state photoconductance data, Applied Physics Letters 69 (1996) 2510–2512.
[46] J.S. Swirhun, R.A. Sinton, M.K. Forsyth, T. Mankad, Contactless measurment of minority carrier lifetime in silicon ingots and bricks, Progress in Photovoltaics: Research and Applications (2010), http://dx.doi.org/10.1002/pip.1029.
[47] K. Bothe, R. Krain, R. Brendel, R. Falster, R. Sinton, Determining the bulk lifetime of unpassivated multicrystalline silicon wafers, in: Proceedings of the Twenty fifth EUPVSEC, 2010, 1828–1833.
[48] J.N. Ximello-Quiebras, J. Junge, S. Seren, G. Hahn, E. Epple, Solutions Used in the Texturization of Monocrystalline Siliconsixteenth ed., Photovoltaics International 68-74.
[49] D.L. King, M.E. Buck, Experimental optimization of an anisotropic etching process for random texturization of silicon solar cells, in: Proceedings of the IEEE PVSC, 1991, pp.303–308.
[50] J. Bultman, I. Cesar, B. Geerligs, Y. Komatsu, W. Sinke, Methods of emitter formation for crystalline silicon solar cells, Photvoltaics International, eighth ed., (2010) 69–80.
[51] A. Hauser, G. Hahn, M. Spiegel, Comparison of different techniques for edge isolation, in: Proceedings of the Seventeenth EU PVSEC, Munich, Germany, 2001, pp. 1739–1742.
[52] E.L. Ralph, Recent advancements in low cost solar cell processing, in: Proceedings of the IEEE PVSC, 1975, Scottsdale, AZ. 315-316.
[53] M.M. Hilali, A. Rohatgi, C. Khadilkar, S. Kim, T. Pham, J. Salami, A. Shaikh, S. Sridharan, Understanding and development of Ag pastes for silicon solar cells with high sheet-resistance emitters, in: Proceedings of the Nineteenth EUPVSEC, 2004, pp. 1300–1303.
[54] R.M. Swanson, Approaching the 29% limit efficiency in silicon solar cells, in: Proceedings of the Twentyth EU PVSEC, Barcelona, Spain, 2005 pp. 584–589.
[55] T. Tiedje, E. Yablonovitch, G.D. Cody, B.G. Brooks, Limiting efficiency of silicon solar cells, IEEE Transactions on Electron Devices ED-31 (1984) 711–716.
[56] J. Nelson, The Physics of Solar Cells, Imperial College Press, London, 2003.
[57] H.J. Queisser, Detailed balance limit for solar cell efficiency, Materials Science and Engineering B 159-160 (2009) 322–328.
[58] D. Kray, M. Hermle, S.W. Glunz, Theory and experiments on the back side reflectance of silicon wafer solar cells, Progress in Photovoltaics: Research and Applications 16 (2008) 1–15.
[59] M. Horteis, S.W. Glunz, Fine line printed silicon solar cells exceeding 20% efficiency, Progress in Photovoltaics: Research and Applications 16 (2008) 555–560.
[60] S. Wenham, Buried-contact silicon solar cells, Progress in Photovoltaics: Research and Applications 1 (1993) 3–10.
[61] D.-H. Neuhaus, A. Munzer, Industrial silicon wafer solar cells, Advances in Optoelectronics (2007) 24521.
[62] R.M. Swanson, Developments in Silicon solar cells, in: IEEE Electron Devices Meeting, 2007, 359–362.
[63] K.R. McIntosh, N.C. Shaw, J.E. Cotter, Light trapping in Sunpower's A-300 solar cells, in: Proceedings of the Nineteenth EUPVSEC, 2004, pp. 844–847.
[64] M.A. Green, Solar Cells: Operating Principles, Technology and System Applications, The University of New South Wales, Sydney, Australia, 1998.
[65] W.P. Mulligan, D.H. Rose, M.J. Cudzinovic, D.M. DeCeuster, K.R. McIntosh, D.D. Smith, R.M. Swanson, Manufacture of Solar Cells with 21% Efficiency, Proceedings of the Nineteenth EUPVSEC, 2004.
[66] Y. Tsunomura, Y. Yoshimine, M. Taguchi, T. Baba, T. Kinoshita, H. Kanno, H. Sakata, E. Maruyama, M. Tanaka, Twenty-two percent efficiency HIT solar cell, Solar Energy Materials and Solar Cells 93 (2009) 670–673.
[67] M. Edwards, Efficiencies of 22% at low cost: the future of mass-produced laser-doped selective emitter solar cells, Photovoltaics International 11th Edition (2011) 56–62.
[68] M. Taguchi, Y. Tsunomura, H. Inoue, S. Taira, T. Nakashima, T. Baba, H. Sakata, E. Maruyama, High-efficiency HIT solar cell on thin (< 100 um) silicon wafer, in: Proceedings of the Twenty fourth EUPVSEC, 2009, pp. 1690–1693.
[69] S. Rein, W. Warta, S.W. Glunz, Investigation of carrier lifetime in p-type Cz-silicon: specific limitations and realistic prediction of cell performance, in: Proceedings of the Twenty eighth IEEE PVSC, Anchorage, AK, 2000, pp. 57–60.
[70] S.W. Glunz, S. Rein, J. Knobloch, W. Wettling, T. Abe, Comparison of boron- and gallium- doped p-type Czochralski silicon for photovoltaic applications, Progress in Photovoltaics: Research and Applications 7 (1999) 463–469.
[71] A. Ebong, M. Hilali, V. Upadhyaya, B. Rounsaville, I. Ebong, A. Rohatgi, High efficiency screen-printed planar solar cells on single crystalline silicon materials, in: Proceedings of the Thirty first IEEE PVSC, 2005, pp. 1173–1176.
[72] T. Dullweber, S. Gatz, T. Falcon, H. Hannebauer, High-efficiency rear-passivated screen-printed silicon solar cells, Photovoltaics International 13 (2011) 77–84.
[73] M.B. Spitzer, S.P. Tobin, C.J. Keavney, High-efficiency ion-implanted silicon solar cells, IEEE Transactions on Electron Devices ED-27 (1984) 792–802.
[74] R.D. Westbrook, R.F. Wood, G.E. Jellison, Glow-discharge-implanted, thermally annealed, oxide-passivated silicon solar cells of 19% efficiency, Applied Physics Letters 50 (1987) 469–471.
[75] T. Janssens, N.E. Posthuma, B.J. Pawlak, E. Rosseel, J. Poortmans, Implantation for an excellent definition of doping profiles in Si solar cells, in: Proceedings of the Twenty fifth EUPVSEC, 2010, pp. 1179–1181.
[76] M. Hermle, F. Granek, O. Schultz-Wittmann, S.W. Glunz, Shading effects in back-junction back-contacted silicon solar cells, in: Proceedings of the Thirty third IEEE PVSC, 2008, pp. 1–4.
[77] S. Kluska, F. Granek, M. Hermle, S.W. Glunz, Loss analysis of high-efficiency back-contact back-junction silicon solar cells, in: Proceedings of the Twenty third EUPVSEC, 2008, pp. 1590–1595.
[78] K.R. McIntosh, M.J. Cudzinovic, D.D. Smith, W.P. Mulligan, R.M. Swanson, The choice of silicon wafer for the production of low-cost rear-contact solar cells, in: Third World Conference on Photovoltaic Energy Conversion; Osaka, Japan, 2003, pp. 971–974.
[79] E. Yablonovitch, T. Gmitter, R.M. Swanson, Y.H. Kwark, A 720 mV open-circuit voltage SiOx:c-Si:SiOx double heterostructure solar cell, Applied Physics Letters 47 (1985) 1211–1213.
[80] E.V. Kerschaver, G. Beaucarne, Back-contact solar cells: a review, Progress in Photovoltaics: Research and Applications 14 (2006) 107–123.
[81] M. Tanaka, M. Taguchi, T. Matsuyama, T. Sawada, S. Tsuda, S. Nakano, H. Hanafusa, Y. Kuwano, Development of new a-Si/ c-Si heterojunction solar cells: ACJ-HIT (artifically constructed junction-heterojunction with intrinsic thin-layer Japanese, Journal of Applied Physics 31 (1992) 3518–3522.
[82] S.W. Glunz, A. Mette, M. Aleman, P.L. Richter, A. Filipovic, G. Willeke, New concepts for the front side metallization of silicon solar cells, in: Proceedings of the Twenty first EU PVSEC, Dresden, Germany, 2006, pp. 746–749.
[83] G. Schubert, F. Huster, P. Fath, Current transport mechanism in printed Ag thick film contacts to an n-type emitter of a crystalline silicon solar cell, in: Proceedings of the Nineteenth EUPVSEC, 2004, Paris, p. 813.
[84] M. Edwards, J. Bocking, J.E. Cotter, N. Bennet, Screen-print selective diffusions for high-efficiency industrial silicon solar cells, Progress in Photovoltaics: Research and Applications 16 (2008) 31–45.
[85] D. Ide, M. Taguchi, Y. Yoshimine, T. Baba, T. Kinoshita, H. Kanno, H. Sakata, E. Maruyama, M. Tanaka, Excellent power-generating properties by using the HIT structure, in: Proceedings of the Thirty third IEEE PVSC, 2008, 1–5.
[86] S.W. Glunz, High efficiency crystalline silicon solar cells, Advances in OptoElectronics, 2007 (2007) 15p., http://dx.doi.org/10.1155/2007/97370.
[87] Z. Shi, S. Wenham, J. Ji, S. Partlin, A. Sugianto, Mass production of the innovative Pluto solar cell technology, Suntech White Paper., Available online at: 〈http://am.suntech-power.com/technology/pluto.html〉, 2011.
[88] S.W. Glunz, J. Benick, D. Biro, M. Bivour, M. Hermle, D. Pysch, M. Rauer, C. Reichel, A. Richter, M. Rudiger, C. Schmiga, D. Suwito, A. Wolf, R. Preu, n-type silicon–enabling efficiencies > 20% in industrial production, in: Proceedings of the IEEE PVSC, 2010, pp. 000050–000056.
[89] Z. Wang, P. Han, H. Lu, H. Qian, L. Chen, Q. Meng, N. Tang, F. Gao, Y. Jia, J. Wu, Y. Fei, W. Wu, H. Zhu, J. Ji, Z. Shi, A. Sugianto, L. Mai, B. Hallam, S. Wenham, Advanced PERC and PERL production cells with 20.3% record efficiency for standard commercial p-type silicon wafers, Progress in Photovoltaics: Research and Applications (2012), http://dx.doi.org/10.1002/pip.2178.
[90] S. McMahon, SunPower starts commercial production of 24%-efficient Maxeon Gen 3 cell, PV-Tech, Available online at: 〈http://www.pv-tech.org/news/sunpower_starts_commercial_production_of_24_efficient_maxeon_gen_3_cell〉, March 29, 2012.
[91] M. Hoffman, S.W. Glunz, R. Preu, G. Willeke, 22% efficient silicon solar cells using amorphous silicon rear side passivation, in: Proceedings of the Twenty first EUPVSEC, Dresden, Germany, 2006, pp. 609–613.
[92] H. Kanno, D. Ide, Y. Tsunomura, S. Taira, T. Baba, Y. Yoshimine, M. Taguchi, T. Kinoshita, H. Sakata, E. Maruyama, Over 22% efficient HIT solar cell, in: Proceedings of the Twenty third EUPVSEC, 2008, pp. 1136–1139.
[93] M.A. Green, K. Emery, Y. Hishikawa, W. Warta, E.D. Dunlop, Solar cell efficiency tables, Progress in Photovoltaics: Research and Applications 20 (2012) 12–20.
[94] T.M. Bruton, N.B. Mason, S. Roberts, Towards 20% efficienct silicon solar cells manufactured at 60 MWp/annum, in: Proceedings of the Third World Conference on Photovoltaic Energy Conversion, 1 (2003) pp. 899–902.
[95] J. Zhao, A. Wang, M.A. Green, 24.5% efficiency silicon PERT cells on MCZ substrates and 24.7% efficiency PERL cells on FZ substrates, Progress in Photovoltaics: Research and Applications 7 (1999) 471–474.
[96] Y. Komatsu, Innovative diffusion process for improved efficiency on industrial solar cells by doping profile manipulation, in: Proceedings of the EUPVSEC, 2009, pp. 1063–1067.
[97] R.R. King, R.A. Sinton, R.M. Swanson, Studies of diffused phosphorous emitters: saturation current, surface recombination velocity, and quantum efficiency, IEEE Transactions on Electron Devices 37 (1990) 365–371.
[98] T.M. Pletzer, H. Windgassen, H. Kurz, Applicability of screen-printable dopant pastes: gettering and selective emitters, Photvoltaics International twelfth ed., (2011) 65–71.
[99] F. Colville, Laser-assisted selective emitters and the role of laser doping, Photovoltaics International fifth ed., (2009) 84–91.
[100] D. Kray, M. Aleman, A. Fell, S. Hopman, K. Mayer, M. Mesec, R. Muller, G.P. Willeke, S.W. Glunz, B. Bitnar, D.H. Neuhaus, R. Ludemann, T. Schlenker, D. Manz, A. Bentzen, E. Sauar, A. Pauchard, B. Richerzhagen, Laser-doped silicon solar cells by laser chemical processing (LCP) exceeding 20% efficiency, in: Proceedings of the Thirty third IEEE PVSC, 2008.
[101] H. Antoniadis, F. Jiang, W. Shan, Y. Liu, All screen printed mass produced silicon ink selective emitter solar cells, in: IEEE Photovoltaics Specialist Conference, 2010, 001193–001196.
[102] A.G. Aberle, Surface passivation of crystalline silicon solar cells: a review, Progress in Photovoltaics: Research and Applications 8 (2000) 473–487.

[103] B. Xu, K. Littau, J. Zesch, D. Fork, Front side metallization of crystalline silicon solar cells using selectively laser drilled contact openings, in: Proceedings of the IEEE PVSC, 2009, pp. 000517–000522.
[104] M. Osborne, IMEC and Kaneka demonstrate first heterojunction cell with copper-electroplated top grid electrode, PV-Tech., November 28, 2011.
[105] M. Osborne, Fraunhofer ISE confirms Schott Solar's copper-based contact cell with 19.7% efficiency, PV-Tech., December 12, 2011.
[106] A.A. Istratov, E.R. Weber, Physics of copper in silicon, Journal of the Electrochemical Society 149 (2002) G21–G30.
[107] M. Horteis, J. Bartsch, S. Binder, A. Filipovic, J. Merkel, V. Radtke, S.W. Glunz, Electrical properties of fine line printed and light-induced plated contacts on silicon solar cells, Progress in Photovoltaics: Research and Applications 18 (2010) 240–248.
[108] A. Mette, C. Schetter, D. Wissen, S. Lust, S.W. Glunz, G. Willeke, Increasing the efficiency of screen-printed silicon solar cells by light-induced silver plating, in: Proceedings of the Fourth World Conference on Photovoltaic Energy Conversion, Waikoloa, Hawaii, 2006, pp. 1056–1059.
[109] W. Shockley, W.T. Read, Statistics of the recombination of holes and electrons, Physics Review 87 (1952) 835.
[110] R.N. Hall, Electron-hole recombination in germanium, Physics Review 87 (1952) 387.
[111] J. Schmidt, A.G. Aberle, Accurate method for the determination of bulk minority-carrier lifetimes of mono and multicrystalline silicon wafers, Journal of Applied Physics 81 (1997) 6186–6199.
[112] S. Dauwe, J. Schmidt, R. Hezel, Very low surface recombination velocities on p- and n-type silicon wafers passivated with hydrogenated amorphous silicon films, in: Proceedings of the Twenty ninth IEEE PVSC, 2002, pp. 1246–1249.
[113] M.M. Hilali, A. Rogatgi, S. Asher, Development of screen-printed silicon solar cells with high fill factors on 100 ohm/ square emitters, IEEE Transactions of Electron Devices 51 (2004) 948–955.
[114] P. Lolgen, C. Leguijt, J.A. Eikelboom, R.A. Steerman, W.C. Sinke, L.A. Verhoef, P.F.A. Alkemade, E. Algra, Aluminum back-surface field doping profiles with surface recombination velocities below 200 cm/s, IEEE Transactions on Electron Devices 40 (1993) 236–242.
[115] V. Meemongkolkiat, K. Nakayashiki, D.S. Kim, R. Kopecek, A. Rohatgi, Factors limiting the formation of uniform and thick aluminum back surface field and its potential, Journal of the Electrochemical Society 153 (2006) G53–G58.
[116] S. Narasimha, A. Rohatgi, A.W. Weeber, An optimized rapid aluminum back surface field technique for silicon solar cells, IEEE Transactions on Electron Devices 46 (1999) 1363–1370.
[117] A.W. Blakers, M.A. Green, 20% Efficiency silicon solar cells, Applied Physics Letters 48 (1986) 215.
[118] R.R. King, R.A. Sinton, R.M. Swanson, Front and back surface fields for point-contact solar cells, in: Proceedings of the Twentieth IEEE PVSC, Las Vegas, NV, 1988, p. 538.
[119] Y. Ma, P. Choulat, G. Agostinelli, J. John, X. Loozen, G. Beaucarne, Industrial-type PERC solar cell: Towards 80 microns thickness in: Proceedings of the Twenty third EU PVSEC, 2008, pp. 1193–1196.
[120] G. Agostinelli, A. Delabie, P. Vitanov, Z. Alexieva, H.F.W. Dekkers, S.D. Wolf, G. Beaucarne, Very low surface recombination velocities on p-type silicon wafers passivated with a dielectric with fixed negative charge, Solar Energy Materials and Solar Cells 90 (2006) 3438.
[121] R. Hezel, K. Jaeger, Low-temperature surface passivation of silicon for solar cells, Journal of the Electrochemical Society 136 (1989) 518.
[122] J. Schmidt, T. Lauinger, A.G. Aberle, R. Hezel, Record low surface recombination velocities on low-resistivity silicon solar cell substrates, Proceedings of the Twenty fifth IEEE PVSC, Washington, DC, 1996, pp. 413–416.
[123] I. Martin, M. Vetter, A. Orpella, J. Puigdollers, A. Cuevas, R. Alcubila, Surface passivation of p-type crystalline Si by plasma enhanced chemical vapor deposited amorphous SiCx:H films, Applied Physics Letters 79 (2001) 2199–2201.
[124] S. Janz, S. Riepe, M. Hofmann, S. Reber, S. Glunz, Phosphorous-doped SiC as an excellent p-type Si surface passivation layer, Applied Physics Letters 88 (2006) 133516.
[125] J. Schmidt, A.G. Aberle, R. Hezel, Investigation of carrier lifetime instabilities in Cz-grown silicon, in: Proceedings of the Twenty sixth IEEE PVSC, Anaheim, CA, 1997, pp. 13–18.
[126] M.J. Kerr, J. Schmidt, A. Cuevas, J.H. Bultman, Surface recombination velocity of phosphorous-diffused silicon solar cell emitters passivated with plasma enhanced chemical vapor deposited silicon nitride and thermal silicon oxide, Journal of Applied Physics 89 (2001) 3821–3826.
[127] J.Y. Lee, Ph.D. Dissertation: Rapid thermal Processing of Silicon Solar Cells—Passivation and Diffusion, in: Fraunhofer Institute for Solar Energy Systems; Heidenhofstr, Freiburg, Germany, 2003.
[128] S.W. Glunz, J. Dicker, D. Kray, J.Y. Lee, R. Preu, S. Rein, E. Schneiderlochner, J. Solter, W. Warta, G. Willeke, High-efficiency cell structures for medium quality silicon, in: Proceedings of the Seventeenth EU PVSEC, 2001, pp. 1287–1292.
[129] P. Ortega, A. Orpella, I. Martin, M. Colina, G. Lopez, C. Voz, M.I. Sanchez, C. Molpeceres, R. Alcubilla, Laser-fired contact optimization in c-Si solar cells, Progress in Photovoltaics: Research and Applications (2011), http://dx.doi.org/10.1002/ pip. 1115.
[130] E. Schneiderlochner, R. Preu, R. Ludemann, S.W. Glunz, Laser-fired rear contacts for crystalline silicon solar cells, Progress in Photovoltaics: Research and Applications 10 (2002) 29–34.
[131] W. Koch, A.L. Endros, D. Franke, C. Habler, J.P. Kalejs, H.J. Moller, Bulk crystal growth and wafering for PV, in: A Luque, S Hegedus (Eds.), Handbook of Photovoltaic Science and Engineering, John Wiley & Sons Ltd., West Sussex, England, 2003.
[132] H. Fischer, W. Pschunder, Investigation of photon and thermal induced changes in silicon solar cells, in: Proceedings of the Tenth IEEE PVSC, 1973, pp. 404–411.
[133] J.E. Cotter, J.H. Guo, P.J. Cousins, M.D. Abbott, F.W. Chen, K.C. Fisher, P-type versus n-type silicon wafers: prospects for high-efficiency commercial silicon solar cells, IEEE Transactions on Electron Devices 53 (2006) 1893–1901.
[134] S.W. Glunz, S. Rein, J. Knobloch, Stable Czochralski silicon solar cells using gallium-doped base material, in: Proceedings of the 16th EU PVSEC, Glasgow, UK, 2000, pp. 1070–1075.
[135] T. Saitoh, H. Hashigami, S. Rein, S.W. Glunz, Overview of light degradation research on crystalline silicon solar cells, Progress in Photovoltaics: Research and Applications 8 (2000) 537–547.
[136] H. Fischer, W. Pschunder, Investigation of photon and thermal induced changes in silicon solar cells, in: Proceedings of the IEEE PVSC, Palo Alto, CA, 1973, pp. 404–411.
[137] J. Knobloch, S.W. Glunz, V. Henninger, 21% efficient solar cells processed from Czochralski grown silicon, in: Proceedings of the Thirteenth EU PVSEC, Nice, France, 1995, pp. 9–12.
[138] S.W. Glunz, S. Rein, W. Warta, J. Knobloch, W. Wettling, On the degradation of Cz-silicon solar cells, in: Proceedings of the Second World Conference on Photovoltaic Energy Conversion, Vienna, Austria, 1998, pp. 1343–1346.
[139] S.W. Glunz, S. Rein, J.Y. Lee, W. Warta, Minority carrier lifetime degradation in boron-doped Czochralski silicon, Journal of Applied Physics 90 (2001) 2397–2404.
[140] J. Schmidt, Light-induced degradation in crystalline silicon solar cells, Solid State Phenomena 95-96 (2004) 187–196.
[141] K. Bothe, R. Sinton, J. Schmidt, Fundamental boron-oxygen-related carrier lifetime limit in mono- and multi-crystalline silicon, Progress in Photovoltaics: Research and Applications 13 (2005) 287–296.
[142] I. Rutschmann, A call for quality: power loss from crystalline silicon module degradation causes a big headache for the industry, Photon International 3 (2008) 106–111.
[143] R.R. King, R.M. Swanson, Studies of diffused boron emitters: saturation current, surface recombination velocity, and quantum efficiency, IEEE Transactions on Electron Devices 38 (1991) 1399–1409.
[144] V. Meemongkolkiat, K. Nakayashiki, A. Rohatgi, G. Crabtree, J. Nickerson, T.L. Jester, Resistivity and lifetime variation along commercially grown Ga- and B-doped Czochralski Si ingots and its effect on light-induced degradation and performance of solar cells, Progress in Photovoltaics: Research and Applications 14 (2006) 125–134.
[145] S.K. Pang, A. Rohatgi, Record high recombination lifetime in oxidized magnetic Czochralski silicon, Applied Physics Letters 59 (1991) 195–197.
[146] S.K. Chunduri, Applying magnetic fields for better crystal quality, Photon International 6 (2010) 207.
[147] A. Cuevas, C. Samundsett, M.J. Kerr, D.H. Macdonald, H. Mackel, P.P. Altermatt, Back junction solar cells on n-type multicrystalline and Cz silicon wafers, in: Proceedings of the Third World Conference on Photovoltaic Energy Conversion, 2003, pp. 963–966.
[148] D. Macdonald, L.J. Geerlings, Recombination activity of interstitial iron and other transition metal point defects in p- and n-type crystalline silicon, Applied Physics Letters 85 (2004) 4061–4063.
[149] F. Granek, M. Hermle, D.M. Huljic, O. Schultz-Wittmann, S.W. Glunz, Enhanced lateral current transport via the front n+ diffused layer of n-type high-efficiency back-junction back-contact silicon solar cells, Progress in Photovoltaics: Research and Applications 17 (2009) 47–56.
[150] V. Mertens, S. Bordihn, Y. Larinova, N.P. Harder, R. Brendel, The buried emitter solar cell concept: Interdigitated back-junction structure with virtually 100% emitter coverage of the cell area, in: Proceedings of the Twenty fourth EUPVSEC, 2009, pp. 934–936.
[151] D. DeCeuster, P. Cousins, D. Rose, D. Vicente, P. Tipones, W. Mulligan, Low cost, high volume production of >22% efficiency silicon solar cells, in: Proceedings of the Twenty second EUPVSEC, 2007, pp. 816–819.
[152] C. Reichel, M. Reusch, F. Granek, M. Hermle, S.W. Glunz, Decoupling charge carrier collection and metallization geometry of back-contacted back-junction silicon solar cells by using insulating thin films, inProceedings of the Thirty sixth IEEE PVSC, 2010, pp. 001034–001038.
[153] Applied materials completes acquisition of Varian Semiconductor Equipment Associates, Applied Materials Press Release. Available online at: 〈http://www.appliedmaterials.com/varianacquisitionclosserelease〉, November 10, 2011.
[154] A. Gupta, R.J. Low, N.P.T. Bateman, D. Ramappa, H.J. Gossman, Q. Zhai, P. Sullivan, W. Skinner, C. Dube, B. Tsefrekas, J. Mullin, High efficiency selective emitter cells using in-situ patterned ion implantation, in: Proceedings of the Twenty fifth EUPVSEC, 2010, pp. 1158–1162.
[155] M. Schmela, Solar module survey 2012, Photon International 2 (2012) 130–171.
[156] M. Taguchi, K. Kawamoto, S. Tsuge, T. Baba, H. Sakata, M. Morizane, K. Uchihashi, N. Nakamura, S. Kiyama, O. Oota, HIT cells—high efficiency crystalline Si cells with novel structure, Progress in Photovoltaics: Research and Applications 8 (2000) 503–513.
[157] E. Maruyama, A. Terakawa, M. Taguchi, Y. Yoshimine, D. Ide, T. Baba, M. Shima, H. Sakata, M. Tanaka, Sanyo's challenges to the development of high-efficiency HIT solar cells and the expansion of the HIT business,

Fourth World Conference on Photovoltaic Energy Conversion, 2006, pp. 1455–1460.
[158] H. Fujiwara, M. Kondo, Effects of a-Si:H layer thickness on the performance of a-Si:H/ c-Si heterojunction solar cells, Journal of Applied Physics 101 (2007) 054516.
[159] R.A. Street, D.K. Biegelsen, J.C. Knights, Defect states in doped and compensated a-Si:H, Physical Review B 24 (1981) 969–984.
[160] S. Taira, Y. Yoshimine, T. Baba, M. Taguchi, H. Kanno, T. Kinoshita, H. Sakata, E. Maruyama, M. Tanaka, Our approaches for achieving HIT solar cells with more than 23% efficiency, in: Proceedings of the Twenty second EUPVSEC, 2007, pp. 932–935.
[161] M. Edwards, S. Bowden, U. Das, M. Burrows, Effect of texturing and surface preparation on lifetime and cell performance in heterojunction silicon solar cells, Solar Energy Materials and Solar Cells 92 (2008) 1373–1377.
[162] I. Korte, E. Conrad, H. Angermann, R. Stangl, M. Schmidt, Overview on a-Si:H/ c-Si heterojunction solar cells—physics and technology, in: Proceedings of the Twenty second EUPVSEC, 2007, pp. 859–865.
[163] M. Taguchi, A. Terakawa, E. Maruyama, M. Tanaka, Obtaining a higher Voc in HIT cells, Progress in Photovoltaics: Research and Applications 13 (2005) 481–488.
[164] J.L. Hernandez, K. Yoshikawa, A. Feltrin, N. Menou, N. Valckx, E.V. Assche, J. Poortmans, D. Adachi, M. Yoshimi, T. Uto, H. Uzu, T. Kuchiyama, C. Allebe, N. Nakanishi, T. Terashita, T. Fujimoto, G. Koizumi, K. Yamamoto, High efficiency silver-free heterojunction silicon solar cell, in: Proceedings of the Twenty first International Photovoltaic Science and Engineering Conference, Fukuoka, Japan, 2011.
[165] G. Crabtree, T.L. Jester, C. Fredric, J. Nickerson, V. Meemongkolkiat, A. Rohatgi, Production viability of gallium doped mono-crystalline silicon solar cells, in: Proceedings of the 31st IEEE PVSC, 2005, pp. 935–938.
[166] D.M. Huljic, T. Zerres, A. Mohr, K.v. Maydell, K. Petter, J.W. Muller, H. Feist, N.P. Harder, P. Engelhart, T. Brendemuhl, R. Grischke, R. Meyer, R. Brendel, F. Granek, A. Grohe, M. Hermle, O. Schultz, S.W. Glunz, Development of a 21% back contact monocrystalline silicon solar cell for large scale production, in: Proceedings of the Twenty first EUPVSEC, 2006, pp. 765–768.
[167] S.K. Chunduri, Market survey on crystal growth equipment, Photon International 6 (2010) 184–209.
[168] 〈http://www2.dupont.com/Photovoltaics/en_US/products_services/metallization/solamet_lowtemp_silver.html〉, (Last accessed November 6, 2012).
[169] M. Berger, Y. Liu, International Technology Roadmap for Photovoltaics (ITRPV): 2011 Results, SEMI PV Group Europe, Available online at: 〈http://www.pvgroup.org〉, 2012.
[170] M. Thirsk, Progress in cost reduction in silver pastes for crystalline silicon cells, Photovoltaics International (2011) 41–44.
[171] Solar cell bypass diodes in Silicon Crystalline Photovoltaic Panels, Vishay General Semiconductor White Paper. Available online at: 〈http://www.vishay.com/docs/89398/solarcell.pdf〉, 2011, 1–5.
[172] H. Wirth, U. Eitner, Current status and future potential of back-contact (BC) module technology14th Edition, Photovoltaics International 122-129.
[173] J. Lillian, The soldering alternative: are conductive adhesives worth another look? Solar Industry 5 (2012) 24–25.
[174] R. Brendel, Thin-film crystalline silicon mini-modules using porous Si for layer transfer, Solar Energy 77 (2004) 969–982.
[175] B. Lalaguna, P. Sanchez-Friera, I.J. Bennett, D. Sanchez, L.J. Caballero, J. Alonso, Evaluation of bismuth-based solder alloys for low-stress interconnection of industrial crystalline silicon PV cells, in: Proceedings of the Twenty seconnd EUPVSEC, 2007, pp. 2712–2715.
[176] A. Goodrich, T. James, M. Woodhouse, Residential, Commercial, and Utility-Scale Photovoltaic (PV) System Prices in the United States: Current Drivers and Cost-Reduction Opportunities, Available online at: 〈http://www.nrel.gov/docs/fy12osti/53347.pdf〉, 2012.
[177] U. S. Department of Energy's Sunshot Program,, Available online at: 〈http://www1.eere.energy.gov/solar/sunshot/〉,.