

PAPER • OPEN ACCESS

Guaranteed global optimization of thin-film optical systems

To cite this article: Paul Azunre *et al* 2019 *New J. Phys.* **21** 073050

View the [article online](#) for updates and enhancements.



PAPER

Guaranteed global optimization of thin-film optical systems

OPEN ACCESS

RECEIVED
26 March 2019REVISED
24 June 2019ACCEPTED FOR PUBLICATION
1 July 2019PUBLISHED
24 July 2019

Original content from this
work may be used under
the terms of the [Creative
Commons Attribution 3.0
licence](#).

Any further distribution of
this work must maintain
attribution to the
author(s) and the title of
the work, journal citation
and DOI.

Paul Azunre^{1,2}, Joel Jean¹, Carmel Rotschild³ , Vladimir Bulovic¹, Steven G Johnson⁴ and Marc A Baldo¹¹ Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of America² Algorine Inc., Austin, TX 78729, United States of America³ Department of Mechanical Engineering, Technion Institute of Technology, Technion City, Haifa 32000, Israel⁴ Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139, United States of AmericaE-mail: azunre@algorine.com**Keywords:** deterministic global optimization, thin film optical filters, solar energy, optical system design**Abstract**

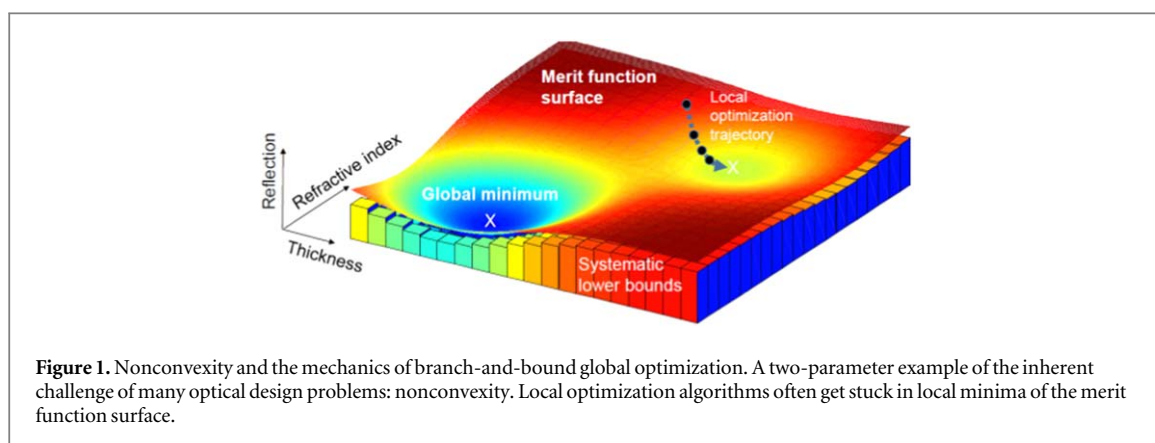
A parallel deterministic global optimization algorithm for thin-film multilayer optical coatings is developed. This algorithm enables locating a global solution to an optimization problem in this class to within a user-specified tolerance. More specifically, the algorithm is a parallel branch-and-bound method with applicable bounds on the merit function computed using Taylor models. This study is the first one, to the best of our knowledge, to attempt guaranteed global optimization of this important class of problems, thereby providing an overview and an assessment of the current state of such techniques in this domain. As a proof of concept on a small scale, the method is illustrated numerically and experimentally in the context of antireflection coatings for silicon solar cells—we design and fabricate a three-layer dielectric stack on silicon that exhibits an average reflectance of $(2.53 \pm 0.10)\%$, weighted over a broad range of incident angles and the solar spectrum. The practicality of our approach is assessed by comparing its computational cost relative to traditional stochastic global optimization techniques which provide no guarantees on their solutions. While our method is observed to be significantly more computationally expensive, we demonstrate via our proof of concept that it is already feasible to optimize sufficiently simple practical problems at a reasonable cost, given the current accessibility of cloud computing resources. Ongoing advances in distributed computing are likely to bring more design problems within the reach of deterministic global optimization methods, yielding rigorous guaranteed solutions in the presence of practical manufacturing constraints.

1. Introduction

Multilayer filters are an integral component of modern optical systems. They function to determine the spectral composition and intensity of light reflected and/or transmitted by an optical system. Perhaps the most important implementation of multilayer filters uses thin films, generally considered to vary between a fraction of a nanometer (nm) and a couple micrometers in thickness [1, 2].

The variety of materials that are available for thin-film deposition and the nm-level control over their thickness lead to a corresponding practical design challenge: what is the *best* multilayer design for a given reflection/transmission spectrum with a given menu of materials? Unfortunately, as illustrated schematically in figure 1, the problem is ‘nonconvex’ on the search space, meaning that multiple suboptimal local minima may exist, potentially distracting the designer from the global optimum solution.

Although many design techniques have been used to attack this problem [1, 3–8], most prominently the needle optimization technique, all previous approaches have a common feature: they provide no theoretical guarantees that the *global* optimum [9] design has been located. Indeed, even if the designer does correctly identify the global solution, present approaches cannot guarantee that the solution is optimal and that a superior solution does not exist. Even the ostensibly ‘global’ optimization algorithms (such as genetic algorithms [3] or other approaches [10]) that have been applied to the thin-film design problem [1] are stochastic (randomized)



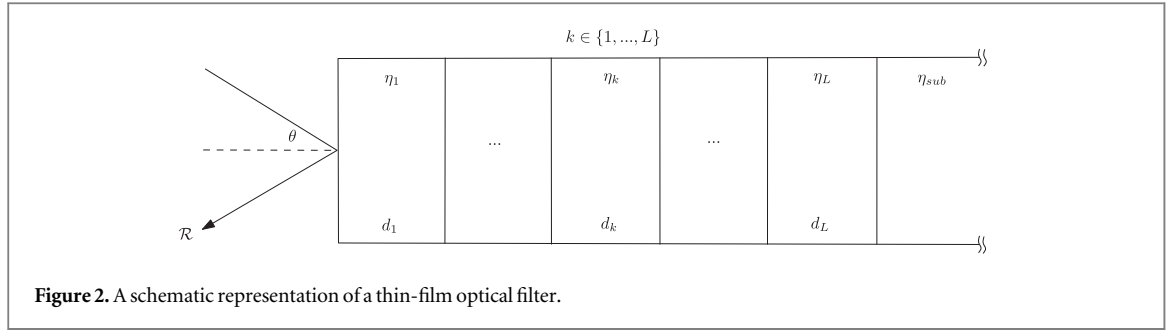
algorithms that converge only *asymptotically* to the global optimum, but provide no information about whether the global optimum has actually been *reached* in the finite number of iterations that is used in practice.

More specifically, approaches to solving nonconvex optics optimization problems are broadly considered to fall into two categories. Local optimization techniques that iteratively improve on an initial guess are known as ‘refinement’ methods [1, 11]. These methods rely on intuition since the initial guess can strongly affect the quality of the final solution. Global optimization methods attempt to overcome nonconvexity and can be categorized as stochastic or deterministic [3]. Stochastic global optimization algorithms escape local minima in some probabilistic manner. Genetic algorithms [12] and other randomized global-search algorithms [13] have been used extensively to design thin-film optical filters. However, convergence of this class of algorithms to the global solution is only guaranteed asymptotically. In practice, execution is terminated when a specified number of iterations or function evaluations has been reached or when the user decides that the current solution is ‘good enough’. Thus, although stochastic algorithms can often avoid suboptimal local minima, one can never know for sure whether a global solution has been found in practice.

Deterministic global optimization approaches, on the other hand, must confront the daunting computational demands of systematically scanning an entire design space. In return, they can provide a guarantee that a globally optimal solution to an optimization problem has been found to within a user-specified tolerance. To determine the practicality of deterministic global optimization, we introduce an algorithm that combines branch-and-bound techniques [9] with calculations of analytic bounds using Taylor arithmetic [14, 15]. We apply this deterministic global algorithm to stochastically-derived solutions for antireflection coatings consisting of two or three layers on a substrate. The deterministic technique is capable of proving that the stochastically-derived solutions are in fact global optima to within a specified tolerance for the specified design constraints in the problems we consider.

As a representative and practical thin-film optics problem, we optimize a broadband omnidirectional antireflection (AR) coating for crystalline silicon (c-Si) solar cells. Existing solutions for reducing reflection losses include surface texturing [16–20], used widely in commercial c-Si cells, in conjunction with a single-layer SiNx:H AR coatings, and gradient-index coatings [21, 22], which approximate a smooth refractive index gradient from that of air to that of silicon. Surface texturing approaches employing nanocone gratings/arrays [23–25], and biomimetic nanostructures [26], have also been explored. Arguably, however, all of these approaches have some serious limitations. Surface texturing can be simple and low-cost, but is generally ineffective for multicrystalline silicon (mc-Si) cells, which currently command over 60% of the c-Si PV market [27], although it has been observed that encapsulation with ethylene vinyl acetate (EVA) reduces the difference in antireflection performance between textured mc-Si and single-crystalline silicon (sc-Si). Gradient-index coatings work well for both sc-Si and mc-Si cells. However, to obtain refractive indices near unity, these coatings often employ porous SiO₂ films that are subject to fouling [22]. Although a continuously-graded index with an infinite number of materials eliminates reflections for omnidirectional incidence over a broad wavelength range [21], it is unclear whether examples of discretized gradient-index approaches achieve the optimal result for the particular material systems and layer structures they consider. In this work, we apply deterministic global optimization to find a guaranteed globally optimal antireflection coating for silicon substrates, based on common and durable materials deposited using industrially proven and scalable coating processes.

As one would expect, we find that the computational cost of proving that a design is a global solution increases dramatically as the design space increases, with the consequence that all but the simplest practical design problems are presently numerically infeasible to attempt using this technique. There are several reasons, however, to suspect that deterministic global optimization techniques will not remain limited to the simplest problems. First, we show that is possible to exploit prior knowledge of problem structure to constrain the design



space in particular scenarios. For example, much more complex antireflection coatings can be deterministically optimized by restricting the design of antireflection coatings to gradient index structures. Second, if present-day trends hold, and computational power continues to increase dramatically, we can expect deterministic global optimization approaches to reach more deeply into the design space of thin film coatings that are economically viable in production. Thus, while conventional stochastic or asymptotic approaches are likely to continue to dominate the initial stages of multilayer filter design, it may be practical in some cases to subsequently perform deterministic global optimization to verify these solutions, or identify superior designs within a given parameter space.

To summarize, this study is the first one to attempt guaranteed global optimization of this important class of problems, thereby providing an overview and an assessment of the current state of such techniques in this domain. It highlights some of the most important numerical and theoretical issues that need to be addressed in order to make deterministic techniques even more practical. We also demonstrate that it is already numerically feasible and practical to solve some simple yet important practical problems using deterministic global optimization techniques. For these kinds of problems, our algorithm can be a useful tool that can be used to verify solutions produced by computationally cheaper algorithms.

The rest of this paper is organized as follows. In section 2 below, we define the design problem that we tackle using deterministic global optimization. The algorithm and the bounding procedure are described in section 3. We then turn to the AR coating application. Section 4 explains how imposing the special gradient-index structure on the solution can be leveraged to make the algorithm more computationally efficient for certain AR coating design problems. Numerical results for the AR coating problem, without imposing any such special structure, are presented in section 5. Section 6 presents the experimental AR results and compares them to the numerical results. The numerical and experimental results are discussed and the paper is concluded in section 7.

2. The multilayer optical design problem

A schematic representation of a multilayer filter is shown in figure 2. In the figure, \mathcal{R} represents the front-interface reflectance. This is the ratio of reflected to incident intensity at any given incident wavelength λ , incident angle θ and polarization s or p, while d_k and η_k respectively denote the physical thickness and the complex refractive index of the k th layer. The complex refractive index of the k th layer is explicitly written as


$$\eta_k = n_k + j\kappa_k, \quad (1)$$

with n_k and κ_k respectively denoting the real and imaginary parts of the refractive index of the k th layer. In general η_k is wavelength dependent, but all problems (with the important exception of experimental re-optimization) considered in this work are nondispersive so that η_k is not a function of wavelength. Moreover, all examples considered feature approximately nonabsorbing materials, i.e. $\kappa_k \approx 0 \forall k$. Finally, the materials are assumed to be isotropic. The algorithm generalizes to absorptive materials directly, and it is straightforward to extend the calculation to dispersive and anisotropic refractive indices by appropriate parameterization with respect to wavelength, and polarization.

The task of optimizing a thin-film optical filter generally involves finding a system configuration that most closely approximates the desired response/reflectance over the relevant bandwidth, incident angle range and specified polarization. For a fixed number of layers L , this task may be posed as

$$\mathbf{p}_{\text{opt}} = \operatorname{argmin}_{\mathbf{p} \in P} (O(\{\{\mathcal{R}_p(\mathbf{p}, \lambda_i, \theta_j), \mathcal{R}_s(\mathbf{p}, \lambda_i, \theta_j)\}_{i=1}^m\}_{j=1}^n, L)). \quad (2)$$

Here, the design variable (or parameter) vector \mathbf{p} is specified by some subset of $\{d_k, \eta_k\}_{k=1}^L$ (i.e. by some subset of the thicknesses and complex refractive indices of the layers in the stack) and the subscripts of \mathcal{R} denote the polarization of incident light. The cost function O measures how closely a given configuration \mathbf{p} approximates the ideal response and perhaps penalizes more complex designs. It is usually a numerical approximation for a



definite integral over specified wavelength and incident angle ranges. The objective function O should be increasing in the number of layers L to reflect the fact that a simpler system is preferable.

If O does not explicitly penalize a more complex system through dependence on L , one can do better and achieve a lower global solution $O(\{\{R_p(\mathbf{p}_{\text{opt}}, \lambda_i, \theta_j), R_s(\mathbf{p}_{\text{opt}}, \lambda_i, \theta_j)\}_{i=1}^m\}_{j=1}^n)$ by increasing the number of degrees of freedom through increasing L , and thereby the dimension of the search space P . When using the algorithm outlined in this paper to address this case, one should start with as few layers as is reasonable and repeat the optimization process for incrementally more layers until an acceptable performance is achieved or until adding more layers does not appreciably improve performance. Hence, we define the global solution to be the solution corresponding to the number of layers immediately after diminishing returns set in.

Technically, this is a Mixed Integer Nonlinear Program since O is nonlinear in \mathbf{p} and the entries of \mathbf{p} may be continuous within an interval or be restricted to a finite number of choices, as in the case where some library of materials is available in the laboratory. However, in this work attention is henceforth restricted to continuous problems.

3. Deterministic global optimization algorithm

In this section, the branch-and-bound algorithm is presented. This framework for deterministic global optimization was first proposed in 1960 by Land and Doig [9]. The algorithm systematically divides up the search space, establishing rigorous lower bounds on each region, and converges when the best candidate solution is sufficiently close to the global lower bound. First, we outline the branching method for systematically dividing up the parameter space, which employs best-bound subinterval and relative-width bisection direction selection rules, and uses midpoint evaluation for candidate-solution search. Secondly, we describe a procedure for bounding functions of the front-interface reflection. The interval bounds are constructed using Taylor arithmetic [14, 15].

3.1. Branch-and-bound method

Here, we assume that the search space P is a closed and bounded interval. Branch-and-bound begins with a crude division of the search space. The objective lower bound on each subinterval/subspace is evaluated following the methods described in the next section. These lower bounds are compared to the presently best candidate solution, which we initially determine using a stochastic optimization algorithm, e.g. genetic algorithms [3, 12] or variants of random multistart [10]. If the lower bound of a particular subspace is larger than the solution then that subspace can be eliminated. Remaining subspaces are divided, candidate solutions computed on each subspace by evaluating the merit function at the midpoint of the interval, and the procedure is repeated until the global solution is identified. Although procedures such as this one are generally considered to be impractically expensive computationally, various stages of it admit readily to parallelization. Given today's availability and low-cost of cloud computing resources, we believe such algorithms will become increasingly important with time. Full details of the branch-and-bound implementation and its parallelization are presented in [28]. Here, we focus on the choice of intervals for bisection.

The bisection of a simple nonconvex univariate objective is illustrated in figure 3. We considered three rules for selecting subintervals for bisection. The first involves picking the first subinterval corresponding to the least remaining lower bound (LRLB). This rule is referred to as the *best-bound* rule in the deterministic global optimization literature. The second involves picking the first subinterval corresponding to the least remaining upper bound. This rule is referred to as the *best-estimate* rule in the literature. The third rule involves picking the first subintervals corresponding to the least bisected subspace. This rule is referred to as the *breadth-first* rule in the literature.

To determine which parameter of the chosen interval to bisect, we defined a quantity

$$\frac{w(I_i)}{w(P_i)} \quad (3)$$

which we term the *relative-width*. The width w is defined as the difference between the upper and lower bounds of the corresponding interval I_i , and i is an index for the parameters in the search space. For example, given a thin film coating of L layers, each described by a refractive index and thickness, $i \in \{1, 2, \dots, 2L - 1, 2L\}$. Normalization by the width of the original space, $w(P_i)$, removes dimensional differences between different parameters.

We found empirically that the best-bound rule coupled with choosing i to maximize the relative-width worked best [28]. The best-bound rule is motivated by the desire to increase LRLB value as fast as possible since this is critical to convergence. Maximizing the relative-width divides the space in a more uniform way.

3.2. Lower bounding procedure

3.2.1. Taylor arithmetic

Once an interval in the design space has been identified, it is necessary to determine the bounds on the objective function within that interval. If the objective function is factorable, i.e. a function that can be computed in a finite number of simple steps, the bounds can be determined using interval arithmetic [29–32]. This is a set of rules corresponding to each step in the computation of the function. For example, consider the function $a + b$. The corresponding interval arithmetic is:

$$[a^L, a^U] + [b^L, b^U] = [a^L + b^L, a^U + b^U], \quad (4)$$

where the square brackets $[\cdot]$ denote an interval bound for the enclosed quantity, and the superscripts L and U signify its lower and upper bounds, respectively. The tightness of the resulting bounds, however, may suffer from the *dependency problem*. Unless the objective function can be arranged such that each interval-valued variable appears only once, interval arithmetic may be unable to account for the dependency (or sensitivity) of each term on the underlying independent variables. For example, consider the function $a - a = 0$. Naïve application of interval arithmetic yields

$$[a^L, a^U] - [a^L, a^U] = [a^L, a^U] + [-a^U, -a^L] = [a^L - a^U, a^U - a^L] \neq 0. \quad (5)$$

The dependency problem is a serious issue for thin-film optical filters. Consider the closed-form expression for reflectance in the relatively simple case of a two layer device at normal incidence, where we assume that all materials are nonabsorbing so that the refractive indices are real, and where n_{sub} stands for the refractive index of the substrate:

$$\mathcal{R}(\mathbf{p}) = \frac{\left((1 - n_{\text{sub}}) \cos \delta_1 \cos \delta_2 - \left(\frac{n_2}{n_1} - n_{\text{sub}} \frac{n_1}{n_2} \right) \sin \delta_1 \sin \delta_2 \right)^2 + \left(\left(\frac{n_{\text{sub}}}{n_1} - n_1 \right) \sin \delta_1 \cos \delta_2 + \left(\frac{n_{\text{sub}}}{n_2} - n_2 \right) \cos \delta_1 \sin \delta_2 \right)^2}{\left((1 + n_{\text{sub}}) \cos \delta_1 \cos \delta_2 - \left(\frac{n_2}{n_1} + n_{\text{sub}} \frac{n_1}{n_2} \right) \sin \delta_1 \sin \delta_2 \right)^2 + \left(\left(\frac{n_{\text{sub}}}{n_1} + n_1 \right) \sin \delta_1 \cos \delta_2 + \left(\frac{n_{\text{sub}}}{n_2} + n_2 \right) \cos \delta_1 \sin \delta_2 \right)^2}. \quad (6)$$

Here,

$$\delta_k = \frac{2\pi n_k d_k}{\lambda} \quad (7)$$

is the phase change experienced by electromagnetic radiation in passing through layer k . Note the multiple occurrences of the refractive index and thickness variables, which, to the best of our knowledge, cannot be eliminated by simply rearranging the expression.

To alleviate the dependency problem, one can employ an approach referred to as Taylor arithmetic and automated in the system COSY INFINITY [33] that is based on Fortran 77. Taylor arithmetic applies Taylor's theorem to bound an $o + 1$ times continuously partially differentiable (on the interval under consideration) function f of the interval-valued variable \mathbf{p} by applying interval arithmetic to the following expression:

$$[f](\mathbf{p}) = f(\mathbf{p}_0) + \sum_{i=1}^o \frac{1}{i!} D^i f(\mathbf{p}_0)(\mathbf{p} - \mathbf{p}_0)^i + [r](\mathbf{p}, \mathbf{p}_0), \quad (8)$$

where $D^i f(\mathbf{p}_0)$ is the i th order partial derivative of f at \mathbf{p}_0 . An explicit expression for the remainder is available for such functions, it being the main tool for obtaining $[r]$; see [33] and the references therein. In other words, after

expressing the function as the sum of its Taylor expansion of some specified order o around some reference point \mathbf{p}_0 and a remainder term r , interval arithmetic is applied to that expression.

In an intuitive sense, the reason this works in reducing the dependency problem is that equation (8) attempts to explicitly express the function in terms of its dependencies (the derivatives, at \mathbf{p}_0) on \mathbf{p} . Derivatives are computed using automatic differentiation [32]. In general, higher o yields tighter bounds at a higher computational cost and is chosen empirically. We choose o arbitrarily in this work, as reported in section 5. The reference point \mathbf{p}_0 may also be chosen arbitrarily. In this work, we choose the midpoint of the interval. Beyond simple application of interval arithmetic to (8), more intelligent use of the Taylor expansion can lead to tighter bounds. In this work, we employ the linear dominated bound as such an intelligent alternative; see [33].

We must now check whether differentiability requirements are satisfied. These requirements are satisfied by merit functions which are smooth functions of \mathcal{R} (of class C^∞) for any o , since the composition of smooth functions is a smooth function, and both the numerator and the denominator of \mathcal{R} are built up entirely from smooth functions of \mathbf{p} (polynomial functions, the trigonometric functions and roots) and binary operations which preserve smoothness. Since both terms are positive, the lower bound on the result of the division is obtained as the division of the lower bound on the numerator and the upper bound of the denominator. The upper bound on the result of the division is obtained as the division of the upper bound on the numerator and the lower bound of the denominator. This can then be used as an interval to construct a bound on the merit function. We note that this is not necessary for the examples we look at in this work, since both involve minimizing reflection, so that only the lower bound of this interval is required. Finally, note that when minimizing the square of reflectance, as one would do when designing a beam splitter for instance, the smoothness requirement is trivially satisfied by observing that the square is a smooth function.

4. Exploiting problem structure: domain reduction using gradient index constraint

In their naïve form, branch-and-bound algorithms scale poorly, and exponentially in the worst case. This motivates the discovery of structural features that may be exploited to speed them up. In this section, we describe one such feature relevant to broadband antireflection coatings, such as the ones that will be designed later.

Dating to Rayleigh's analogy to light propagating without reflection from space into successively denser layers of the atmosphere, the gradient index form has been previously conjectured to be the general solution provided that the bandwidth is sufficiently wide [34, 35]. Indeed, later on we will confirm that the global solutions have the gradient index form: the refractive indices increase monotonically and the thickness of each layer decreases monotonically from air to substrate.

Here, we exploit this form to make the algorithm more efficient. For every variable other than those in the first layer, we create a variable to lie on the interval $[0, 1]$. We call these $\forall i > 1, \gamma_i^{n_{gi}}$. Then, $\forall i > 1$, set $n_{i+1} = \gamma_i^{n_{gi}} n_i$. An analogous constraint can be implemented for the thickness variables. In doing this, the domain is reduced to a fraction that is $\frac{1}{2^{2(\sum_{i=1}^L (i-1))}}$ of the original space. Here, the factor of 2 in the exponent of the denominator is due to the fact that such conditions hold for both the thickness and refractive index variables, while the summation term in the exponent of the denominator is due to the fact that for every inequality in the monotonicity constraint, the space is reduced to half of its original size. This fraction is equal to 1, $\frac{1}{4}$ and $\frac{1}{64}$ for one, two and three layer problems respectively, promising significant reduction in convergence time.

The refractive index fraction is shown for the three layer case in figure 4. The inequalities in question here are


$$n_1 \leq n_2 \leq n_3. \quad (9)$$

Three planes can be seen in the figure, one for each inequality between the refractive index variables. Hence the reduction is $\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} = \frac{1}{8}$ due to the refractive indices and thereby $\frac{1}{64}$ overall.

Such exponential domain reduction promises very efficient algorithms for gradient-index systems. Consider, also, that there are other classes of optical systems such as gradient-index fibers and gradient-index lenses where this constraint is imposed *a priori* for practical manufacturing reasons, so that this domain reduction mechanism may be more widely applicable than thin-film antireflection coatings. We also note that the domain reduction can be used by any optimization algorithm, not just a branch-and-bound or even a deterministic one. Finally, we emphasize that the discussion in this section was purely theoretical, and that in the next section no such special structure is imposed when optimizing.

5. Numerical examples

In this section, we first tackle the particular problem of minimizing average reflection from silicon over the incident angle range $[0^\circ, 60^\circ]$ and the wavelength range $[400, 1600]$ nm with the first example. Although silicon absorption is weak above 1100 nm, reducing reflection at longer wavelengths may be important for



upconversion and silicon photonic applications. We subsequently fine tune the design specifically for silicon solar cells using practical materials, a narrower wavelength range of [400, 1100] nm and the AM1.5G solar spectrum in the second example. Untreated silicon normal incidence reflection is greater than 30% on average over the incident wavelength ranges above, motivating this search for a ‘perfect’ antireflection coating for a silicon solar cell—one that can transmit the most incident light to maximize the efficiency of the solar cell in its relevant wavelength and incident angle ranges of operation, with the ultimate goal of achieving greater efficiency.

We note that although inhomogeneous optical films with continuously variable refractive indices can be fabricated using co-sputtering techniques [36], combinations of discrete layers are typically the easiest to manufacture. Using oblique-angle deposition of SiO₂ and co-sputtering of SiO₂ and TiO₂, it is possible to vary refractive indices continuously in the interval [1.09, 2.60] [36], striking very closely any index needed to approximate a selected profile. This specifies part of our design space.

With regards to prior art, there is a good solution to this ‘perfect’ antireflection coating problem that achieves an average reflection of 3.79%, a seven layer design that approximates the quintic profile

$$n(z) = n_{\min} + (n_{\max} - n_{\min})(10z^3 - 15z^4 + 6z^5). \quad (10)$$

Existing solutions for reducing reflection losses also include surface texturing [16–20, 37], used widely in commercial c-Si cells, in conjunction with a single-layer of SiN_x:H. Unfortunately, surface texturing is generally ineffective for multicrystalline silicon (mc-Si) cells, which currently command over 60% of the c-Si PV market [27].

5.1. Broadband omnidirectional antireflection coating for silicon

The general antireflection design problem is summarized by minimizing the objective

$$O(\mathbf{p}) = \frac{3}{\pi} \frac{1}{1200} \int_0^{\pi/3} \int_{400 \text{ nm}}^{1600 \text{ nm}} \mathcal{R}(\mathbf{p}, \lambda, \theta) d\lambda, \quad (11)$$

where the numerical approximation for the definite integral is performed using the rectangle method with 10 rectangles for each independent variable and the top-left corner approximation, corresponding to $m = 10$ and $n = 10$ in (2). This approximation is also known as the left Riemann sum. We use a 3rd order Taylor expansion (order chosen arbitrarily) for constructing the lower bound on the merit function for this example, as well as for the next example. Thicknesses and refractive indices of every layer are used as design variables thereby specifying the design vector \mathbf{p} . Thicknesses are assumed variable in [5, 500] nm, which we believe to be representative of configurations reliably realizable on our sputtering system, although we do make the thickness interval narrower for the harder problems—as indicated in tables 1 and 2. Refractive indices are assumed to be variable in the interval [1.09, 2.60]. For this merit function, we compare our model to data from the literature [36] and found the discrepancy to be only 0.13% (3.66% versus 3.79% measured in that work). This suggests that our modeling error is less than 0.2% and leads us to set the absolute convergence tolerance to 0.1%. Deterministic algorithm solution information is shown in tables 1 and 2, and stochastic in table 3. Note that the stochastic tool finds the solution—albeit without any guarantee. Also, observe that the three layer solution can be approximated fairly

Table 1. Deterministic solution information for the second numerical example. Wall clock time is presented in the format seconds/hours/days, with some information omitted when redundant. The thickness upper bound d^U is given in the units of nanometers.

L	d^U	Wall clock time	Iterations	Solution (merit, \mathbf{p}_{opt})
1	500	22.5	136	0.112, [1.93, 153]
2	500	118 513/32.9	202 134	0.0526, [1.55, 2.37, 109, 68.3]
2	250	44 093/12.2	110 744	0.0526, [1.55, 2.37, 109, 68.4]
3	200	1663 253/462/19.3	935 599	0.0182, [1.31, 1.85, 2.60, 131, 80.8, 61.9]

Table 2. Deterministic solution cost for the second numerical example. CPU-days here is defined as wall-clock time in days multiplied by the number of processes. The thickness upper bound d^U is given in the units of nanometers. The first L numbers in the parameter vector are refractive indices, and the subsequent L numbers are thicknesses in nm.

L	d^U	CPU-days	EC2 Cost
1	500	0.0042	\$0.30
2	500	21.9	\$6.90
2	250	8.16	\$3.90
3	200	308	\$159

Table 3. Stochastic algorithm solution information for the second design problem.

L	Convergence time (s)	Solution (merit, \mathbf{p}_{opt})
1	66.6	0.112, [1.93 153]
2	295	0.0526, [1.55, 2.37, 109, 68.4]
3	918	0.0182, [1.31, 1.85, 2.60, 131, 81.0, 61.5]

closely using the practical materials MgF_2 , Y_2O_3 and the rutile phase of TiO_2 . MgF_2 could be substituted by other fluorides such as LiF and NaF , Y_2O_3 by other oxides such as HfO_2 and TiO_2 by high index materials such as ZnS .

5.2. Broadband omnidirectional antireflection coating specialized for silicon solar cells

Next, we further refine the design to minimize reflection specifically from silicon solar cells, using the solar spectrum to weight the reflection cost function. The incident angle range considered was from 0° to 60° away from the normal as before, but the wavelength range is narrowed to [400, 1100] nm because weighted silicon absorption for higher wavelengths is negligible. We constrained the solution to a three-layer stack because increasing the number of layers further yields diminishing marginal gains in performance, as shown in [28], and is likely uneconomic in high-volume manufacturing. Moreover, the component materials are initially assumed to be isotropic, non-absorbing, and dispersionless, with refractive indices between $n = 1.09$ and 2.60, consistent with the nanoporous materials described in [36] and the references therein. Under these constraints, we obtained the initial ‘ideal’ solution shown in table 4, corresponding to a global optimum with an average reflectance of $(1.023 \pm 0.10)\%$. We note that the global optimum is a discrete gradient-index solution, i.e. the indices are monotonically increasing from air to silicon. While this property of the solution has already been widely hypothesized in the literature for sufficiently wide bandwidth (see [34] and [35], previous section), this is arguably the most rigorous evidence for this effect to date, since our conclusion is based on a guaranteed global optimum. We note that for this particular problem, a representative stochastic global optimization procedure (MLSL [10]) was also able to find the global optimum in a much shorter time period—albeit without a guarantee of global optimality, as reported in detail in [28]. Thus, it appears that in this case the value of our algorithm is purely theoretical—in proving that some solution is truly a global optimum rather than discovering it. This is not necessarily true in general for other problems in this class that may be feasibly optimized given the current availability and cost of cloud computing resources, and warrants further investigation.

The three-layer structure was subsequently re-optimized over thicknesses for practical materials with refractive index values closest to the initial solution: MgF_2 , Al_2O_3 , and rutile TiO_2 , respectively, from air to




Figure 5. Inspection of the globally optimal coating. (a) SEM cross-section of optimized three-layer antireflection coating on silicon, with average refractive indices shown in parentheses. (b) Picture of the optimized coating under white light. The coated silicon appears dark to the naked eye from all angles.

Table 4. Theoretical solution and experimental realization of guaranteed globally optimal antireflection coatings for silicon. Average reflection for the ideal 3-layer solution, the constrained 3-layer solution and the experimental 3-layer ARC were computed to be 1.02 ± 0.10 , 2.53 ± 0.10 , and $2.76 \pm 0.10\%$ respectively. Average reflection was computed over the wavelength range $400 \text{ nm} \leq \lambda \leq 1100 \text{ nm}$ and for incident angles between 0° and 60° , except the experimental value, which was computed over the same wavelength range and for incident angles between 20° and 60° , due to instrument constraints. For the constrained solution and the experimental result, wavelengths were weighted using the AM1.5G photon flux spectrum, and incident angles were weighted using the benchmark SOLIS model [38, 39]. The particular form of the SOLIS model used was equation (6) of [38], with both parameters in the equation set at 0.9. Experimental error bars are wavelength-averaged values from repeated normal-incidence measurements.

	n_1	$d_1 \text{ (nm)}$	n_2	$d_2 \text{ (nm)}$	n_3	$d_3 \text{ (nm)}$
Ideal 3-layer solution	1.15	139	1.66	87.3	2.60	56.2
Constrained 3-layer solution	1.38	83.5	1.67	39.8	2.43	51.6
Experimental 3-layer ARC		84.6		39.8		51.3

silicon. Wavelengths were weighted using the AM1.5G photon flux spectrum, and incident angles were weighted using the benchmark SOLIS model [38, 39] between 0° and 60° . SOLIS accounts for the diurnal sinusoidal variation in available solar flux, as well as increased atmospheric attenuation at high zenith angles. To obtain experimental refractive index spectra for this final optimization, individual films of each material were prepared on silicon substrates. MgF_2 was deposited by thermal evaporation, while Al_2O_3 and TiO_2 were deposited by RF sputtering (see appendix B). Complex refractive indices of the fabricated films were characterized using a spectroscopic ellipsometer, yielding average values for n of 2.43, 1.67, and 1.38 for TiO_2 , Al_2O_3 , and MgF_2 , respectively. The dispersion of each material for the wavelength range $400 \text{ nm} \leq \lambda \leq 1100 \text{ nm}$ was included in the model. Re-optimization of the film thicknesses yielded the constrained practical solution shown in table 4, with a weighted average reflectance of $(2.53 \pm 0.10)\%$.

6. Experimental characterization

The full ARC stack was fabricated on a silicon wafer, and the thickness of each layer was measured using ellipsometry and confirmed using stylus profilometry (figure B1). Note that although we did not fabricate a complete solar cell—just its front interface optical part—we did theoretically evaluate improvements to the complete solar cell power conversion efficiency via the procedure described in the next two paragraphs and appendix C. Figure 5 shows a scanning electron micrograph (SEM) cross-section of the complete coating and an optical image taken under white light confirming that the coating looks dark to the naked eye. The measured spectral reflectance for a range of incident angles is shown in figure 6 and at normal incidence in figure 7. Low reflectance is observed across the broad range of angles and wavelengths relevant for solar energy harvesting. This experimental result closely matches both the predicted optimum for this set of materials and the modeled reflectance for the final structure. We note that the rapid rise at low wavelengths contributes little to the weighted average reflectance since the AM1.5G spectrum decreases rapidly in that regime.

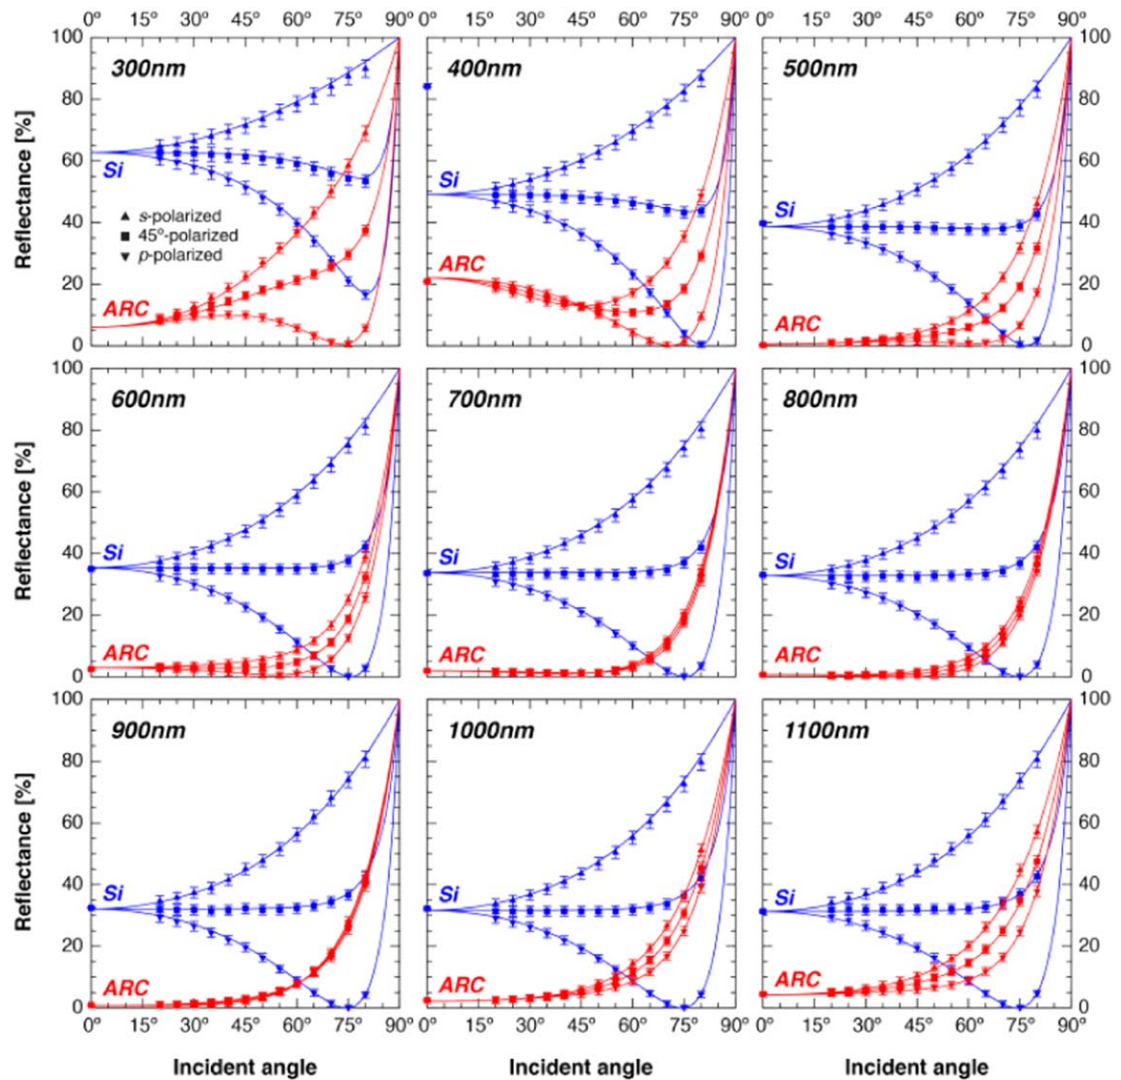


Figure 6. Performance of optimized three-layer antireflection coating. Measured reflectance of bare and AR-coated silicon as a function of incident angle (0° – 80°) away from normal for selected wavelengths and three incident polarizations. Model predictions are shown as solid lines. Low reflectance is observed across the broad range of angles and wavelengths relevant for solar energy harvesting.

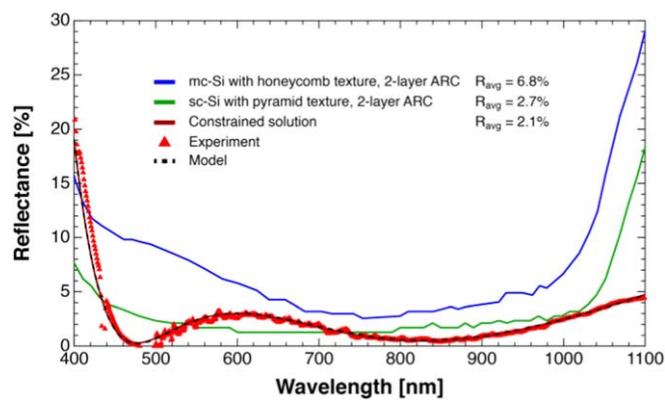


Figure 7. Antireflection performance comparison to high-efficiency c-Si solar cells with and without texturing. Normal-incidence reflectance spectra are shown for our (untextured) three-layer ARC (red), a 24.4%-efficient sc-Si cell with pyramidal texturing (green), and a 19.8% mc-Si cell with honeycomb texturing (blue). AM1.5G-weighted average normal-incidence reflectance values are also shown. Both c-Si cells also employ two-layer (MgF_2/ZnS) antireflection coatings, and reflectance data are adapted from [27]. The theoretically predicted global optimum (dark red), experimental realization (red triangles), and model based on measured thicknesses (black) are shown.

Reflection losses constitute a non-negligible loss pathway in modern mc-Si cells. In figure 7, we compare the normal-incidence reflectance of our optimized coating to that of high-efficiency (among the most efficient to date) sc-Si and mc-Si solar cells from [40]. Both cells employ surface texturing and two-layer antireflection coatings. Our three-layer stack (without any texturing) is observed to outperform these representative devices, with a particularly large difference for the mc-Si cell due to the difficulty of texturing multicrystalline material. However, we note that our use of a 525 μm thick silicon wafer hinders direct comparison with the 260 μm thick back-contacted cells reported by [40]. A discrepancy attributable to parasitic absorption in the back contact and increased optical path length may arise near the band edge, where absorption drops off dramatically and the absorption length starts to exceed the substrate thickness. For a 260 μm thick silicon wafer, nearly all light below 950 nm is absorbed in a single pass. Any underestimation of reflection should thus be limited to wavelengths above 950 nm (see appendix C).

Based on optical performance alone (corrected for wavelengths between 950 and 1200 nm by assuming identical reflectance as the original cell from [40] in that range), application of our optimized three-layer ARC would theoretically increase the AM1.5G power conversion efficiency of the textured mc-Si cell by 0.51% in absolute terms (see appendix C). In practice, the need for passivation of surface and bulk electronic defects may motivate replacement of the bottom TiO_2 layer with high-index $\text{SiN}_x\text{:H}$. These estimates could be considered conservative since they consider normal incidence only—the potential performance gains may be significantly greater at higher angles of incidence [28]. We note that commercial c-Si PV modules employ AR-coated cells encapsulated with EVA ($n = 1.5$) and glass. To enable deployment in a standard module, a re-design or re-optimization of the ARC structure may be required.

7. Conclusion

This paper introduced a deterministic global optimization algorithm for thin-film optical interference coatings, providing an overview and an assessment of the current state of such techniques in this domain. An example problem pertaining to reducing broadband reflection from silicon was studied and characterized experimentally. The practicality of our approach was assessed by comparing its computational cost relative to traditional stochastic global optimization techniques which provide no guarantees on their solutions. While our method is observed to be significantly more computationally expensive, we demonstrate that the current accessibility of cloud computing enables us to verify solutions to simple yet important practical problems. Ongoing advances in distributed computing are likely to bring more design problems within the reach of deterministic global optimization methods, yielding rigorous guaranteed solutions in the presence of practical manufacturing constraints. Moreover, different design problems in this class, of similar complexity but a higher degree of nonconvexity, may benefit from improved solution information as well, not just the guarantee of global optimality.

This work lays the mathematical and algorithmic foundation for the deterministic global optimization of harder optical interference coating design problems, motivating and paving the way for future study of more advanced parallelization techniques that will allow this algorithm to scale to thousands of GPU-enabled CPU nodes so as to enable significantly harder problems to be tackled in the near future. Ultimately we expect that scaleable deterministic global algorithms will transform optical filter design, which is still sometimes regarded as an art, into a science.

Acknowledgments

This material is based upon work supported as part of the MIT Center for Excitonics, an Energy Frontier Research Center funded by the US Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number DE-SC0001088, which sponsored Paul Azunre's doctoral dissertation work. Part of the funding was also provided by the Deshpande Center for Technological Innovation at MIT, which sponsored Paul Azunre's postdoctoral work. Steven G Johnson gratefully acknowledges financial support from the US Army Research Laboratory and the US Army Research Office through the Institute for Soldier Nanotechnologies under contract number W911NF-6 13-D-0001. Joel Jean gratefully acknowledges financial support from the National Science Foundation. The authors would like to thank Nick Thompson, Johanna Engel, Dan Congreve, Markus Einzinger, Matthias Bahlke, Tim McClure, Dr Shiahn Chen, Patrick Boisvert, Elisabeth Shaw, and Kurt Broderick for assistance with the experiments. The authors would also like to thank Collin Fisher Perkinson for his help in preparing the manuscript.

Appendix A. Materials and methods, numerical

In this appendix, we provide some further details pertaining to the numerical implementation of the parallel algorithm that was described in this paper. We first provide more detail on the parallelization and other related issues behind the implementation. This is followed by the numerical exploration of the various properties of the algorithm in the context of a simple and relevant example—one and two layer normal incidence multilayer ARC design problem with carefully selected problem settings. For the most judicious reader, additional details can be found in [28].

A.1. Parallel algorithm implementation issues

The parallel branch-and-bound algorithm was implemented on Amazon's EC2 platform using the COSY INFINITY system [33]. In particular, single work queue dynamic scheduling components of the algorithm on any given server were implemented using the scheduling construct PLOOP made recently available by COSY INFINITY's authors, this construct providing an interface to the Message Passing Interface but restricted to all-to-all communication between processes. It is emphasized that because this construct only allows all-to-all communication between running processes, the communication and synchronization costs prohibit dynamic scheduling of tasks across multiple servers. Further details on this in [28].

All models were tested against analytic examples in the literature while also being validated against real data. Merit lower bounding code was tested for consistency, i.e. we verified that the bounds become tighter as the parameter interval on which the lower bound is computed is made smaller, and that the merit value is attained on a thin/degenerate interval. This allows us to conjecture that our algorithm is provably convergent [9]. Moreover, the model was tested against real data in the context of the numerical example below, and found to be accurate, which validates our modeling assumptions, choice of integral approximation, etc.

Computing costs are reported in CPU-days where appropriate, so as to make it invariant to future computing cost variation. Comparison with stochastic global optimization methods is made where appropriate, the goal of such a comparison being to gauge what advantage this algorithm has over 'state-of-the-art' existing methods. We chose MLSL as a representative stochastic procedure. Reasons for this choice are discussed further in [28]. Termination criteria for MLSL is the maximum number of merit function evaluations and an absolute convergence tolerance on the merit function value, whichever is reached first. The midpoint of the search space is used as an initial guess for the search. In each case, specific values for these termination criteria are given. Optimal system configurations are reported as a vector of length $2L$, with the first L entries representing refractive indices and the rest representing corresponding thicknesses in nm.

When our optimization algorithm was implemented on Amazon's EC2 cloud computing platform, we used a hybrid scheduling approach in the parallelization process: static scheduling at the server level and dynamic scheduling via a single work queue within each server. In particular, convergence for the three-layer omnidirectional problem (that was described in the main text and experimentally demonstrated) was achieved in approximately 3 weeks on a small-scale 16-processor implementation at a cost of approximately \$160 and 309 CPU-days (CPU-days being defined as wall-clock time in days multiplied by the number of processes). See [28] for detailed convergence results and further implementation details.

We next look in detail at a simple numerical example to verify algorithm behavior.

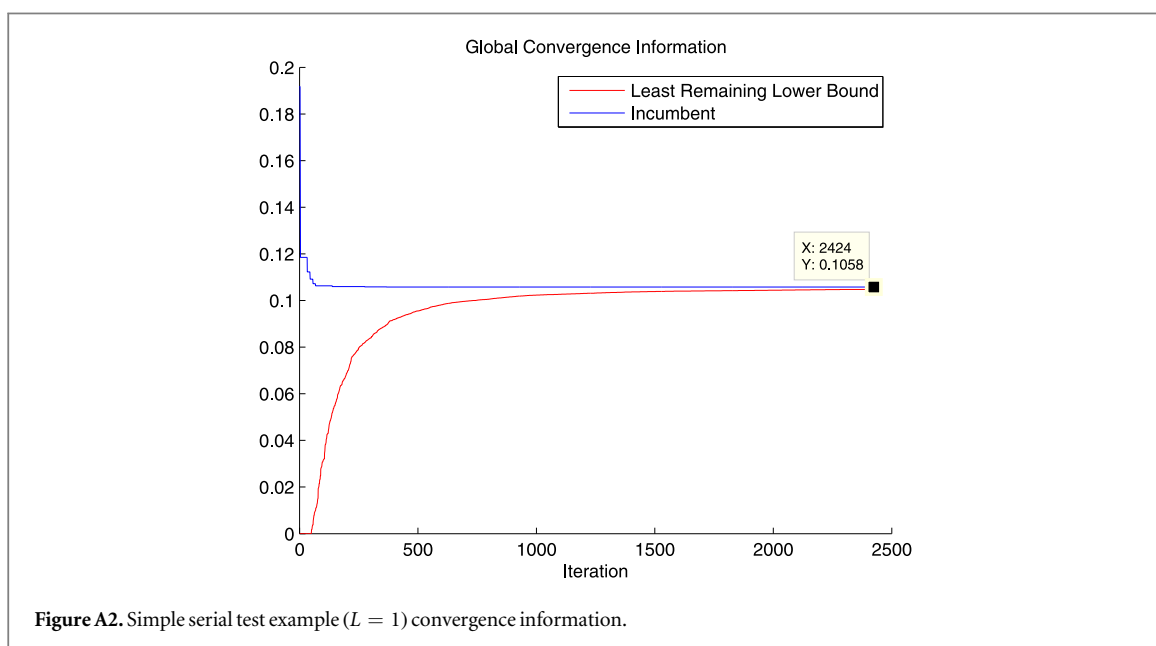

A.2. Optimization of broadband normal incidence antireflection coating for silicon solar cells

Consider the problem of designing an antireflection coating for silicon solar cells, with the goal of minimizing average normal incidence reflectance over a broad range of wavelengths ($\lambda \in [400, 1600]$ nm). This is captured by minimizing the objective

$$\begin{aligned} O(\mathbf{p}) &= \frac{1}{1200} \int_{400 \text{ nm}}^{1600 \text{ nm}} \mathcal{R}(\mathbf{p}, \lambda, 0) d\lambda, \\ &\approx \frac{1}{1200} \frac{1200}{10} \sum_{i=1}^{10} \mathcal{R}(\mathbf{p}, \lambda_i, 0), \quad \lambda_i = 400 + (i-1) \frac{1200}{10} \text{ nm}, \end{aligned} \quad (\text{A.1})$$

where the numerical approximation for the definite integral is performed using the rectangle method (using 10 rectangles and the top-left corner approximation, corresponding to $m = 10$ and $n = 1$ in (2)). We use a 3rd order Taylor expansion (chosen arbitrarily) for constructing the lower bound on the merit function. Thicknesses and refractive indices of every layer are used as design variables (thereby specifying the design vector \mathbf{p}). Thicknesses are assumed variable in $[5, 500]$ nm. Refractive indices are assumed to be variable in the interval $[1.09, 2.60]$. This is consistent with the recent demonstration of refractive index variability achievable through oblique angle deposition of SiO_2 and co-sputtering of $\text{SiO}_2/\text{TiO}_2$ [36]. More details on the model and parameter values in [28].

The point of the numerical exercise in this subsection is to demonstrate that our algorithm is correct for a simple case. We start doing this by thoroughly visualizing its behavior in the context of the one layer (i.e. two



parameter) problem. With only two parameters, it is possible to visualize the merit function and pick the global optimum approximately visually. We show this plot in figure A1. It is clear that this problem is highly nonconvex, even in this small dimensional case, and that the issue can be reasonably expected to become much worse in larger dimensions. We see that the global solution is approximately $\mathbf{p}_{\text{opt}} = [1.95, 145]$, which corresponds to a merit function value of 10.6%. This problem is simple enough for our branch-and-bound algorithm to solve relatively quickly on a single process. Doing this yields the solution $\mathbf{p}_{\text{opt}} = [1.93, 148]$ and a corresponding merit function value of 10.6%, in 2424 iterations, 6.91 s CPU time and 7.07 s wall clock time. This design can be realized approximately experimentally using a material such as yttrium oxide (Y_2O_3). The convergence information (the incumbent and the LRLB evolution) is shown in figure A2. This exercise validates our code.

Table A1. Deterministic algorithm scaling test.

NP	Wall clock (s)	Efficiency	Iterations	Solution (optimal merit, \mathbf{p}_{opt})
1	10 800	1	127 731	0.0465,[1.57, 2.42, 102, 63.0]
4	1410	1.92	31 930	0.0465,[1.57, 2.42, 102, 63.0]
8	608	2.22	15 960	0.0465,[1.57, 2.42, 102, 63.0]
12	384	2.35	10 634	0.0465,[1.57, 2.42, 102, 63.0]
16	302	2.24	7969	0.0465,[1.57, 2.42, 102, 63.0]
32	187	1.81	5717	0.0465,[1.57, 2.42, 102, 63.0]

Next, we increase the complexity of the problem to two layers (i.e. 4 parameters) but increase the absolute convergence tolerance to 2%. This yields a problem that is complex enough to analyze for scaling with number of processes, but simple enough for this to be done in reasonable time. Results of scaling tests are reported in table A1. Efficiency of parallelization is measured as

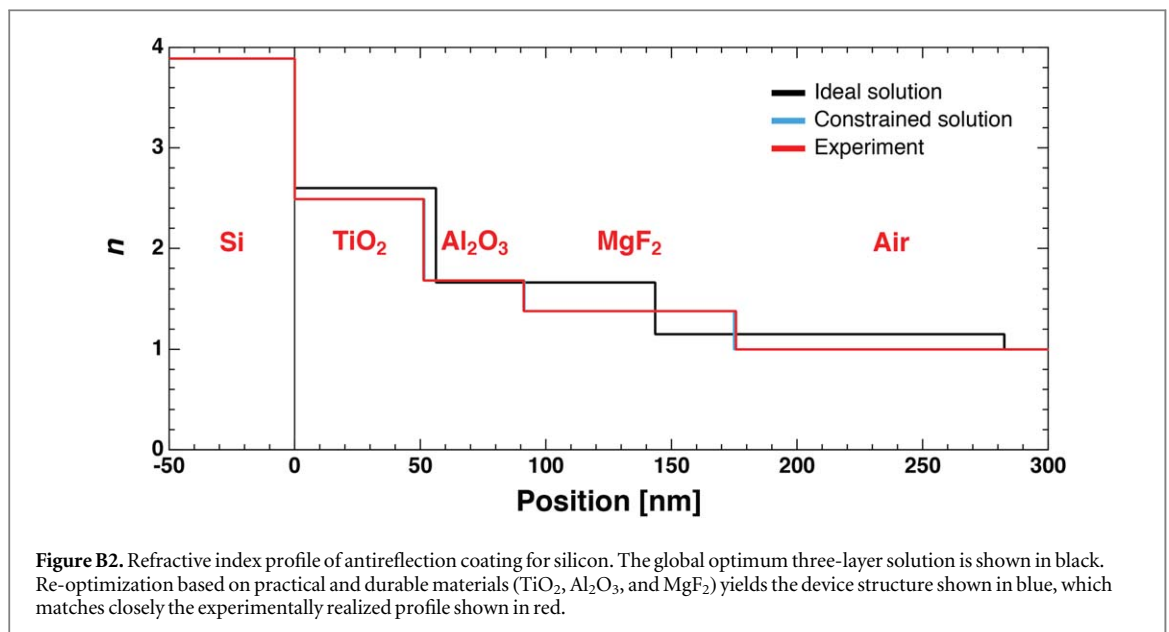

$$E = \frac{T_1}{NPT_{NP}}. \quad (\text{A.2})$$

Here, T_1 is the serial execution time and T_{NP} is the time for execution on NP processors. It is well-established that ideal (linear) scalability would be represented by efficiency of $1 \forall NP$, but many practical algorithms show an efficiency that declines with larger NP due to more effort spent on synchronization and communication (with efficiency reaching 0 for an infinite number of processors) [41]. Efficiency numbers greater than one indicate superlinear speedup. More discussion of these effects in [28].

Appendix B. Materials and methods, experimental

Multilayer antireflection coatings were fabricated on p-type [100] silicon wafers with a thickness of 525 μm . TiO_2 and Al_2O_3 films were deposited by RF sputtering from 3 inch diameter targets at a rate of 0.2–0.4 \AA s^{-1} , Ar pressure of 3 mTorr, and RF power of 200 W. MgF_2 was thermally evaporated from a tungsten boat at 1 \AA s^{-1} at a base pressure of 10^{-6} Torr. No substrate heating was used. Refractive index and extinction coefficient spectra were extracted by fitting data obtained with a Woollam variable-angle spectroscopic ellipsometer. Refractive indices and extinction coefficients for the materials used are shown in figure B1. Average refractive index profiles for the ARC structure are shown in figure B2. Angle-dependent reflectance spectra were measured with the same Woollam instrument and verified with a Cary 500i spectrophotometer equipped with a Variable Angle Spectral Reflectance Accessory. Normal-incidence reflectance measurements were performed using Filmetrics F20 and F40 reflectometers. Corresponding normal-incidence spectra are shown in figure B3, and angle-dependent spectra visualized in figures B4 and B5. SEM imaging was performed using an FEI Helios NanoLab 600i in immersion mode at 5 kV. Sample cross-sections were prepared by focused ion beam milling at 30 kV on the same instrument.

As discussed in the main text, our use of a 525 μm thick silicon wafer may hinder direct comparison of antireflection performance to the 260 μm thick back-contacted cells reported in Zhao *et al* [40]. A thinner wafer absorbs less light, while a metallic back contact increases parasitic absorption but also increases reflection at longer wavelengths where silicon is weakly absorbing. These effects should affect our EQE estimates only for wavelengths above 950 nm, at which more than 1% of light remains unabsorbed after propagating through 260 μm of silicon. To avoid overestimating ARC performance, we calculate theoretical EQE spectra assuming the same reflectance as the original cells for wavelengths between 950 and 1200 nm.



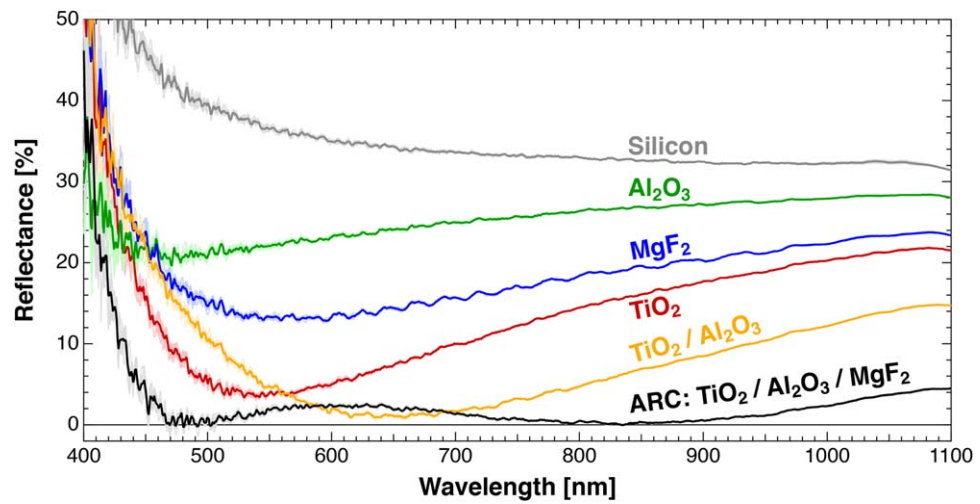


Figure B3. Normal-incidence reflectance for individual films and ARC substacks on silicon.

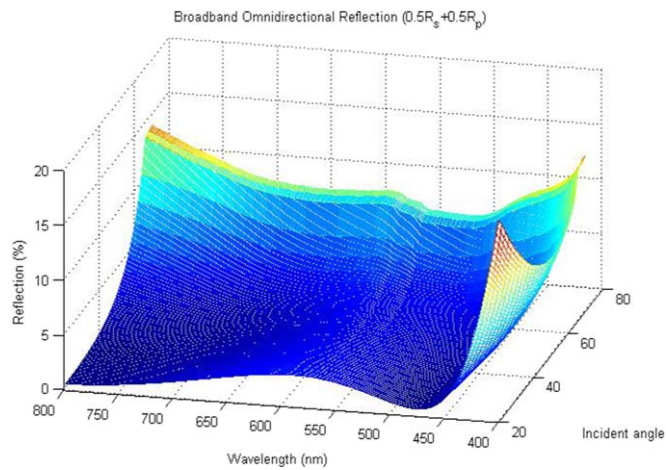
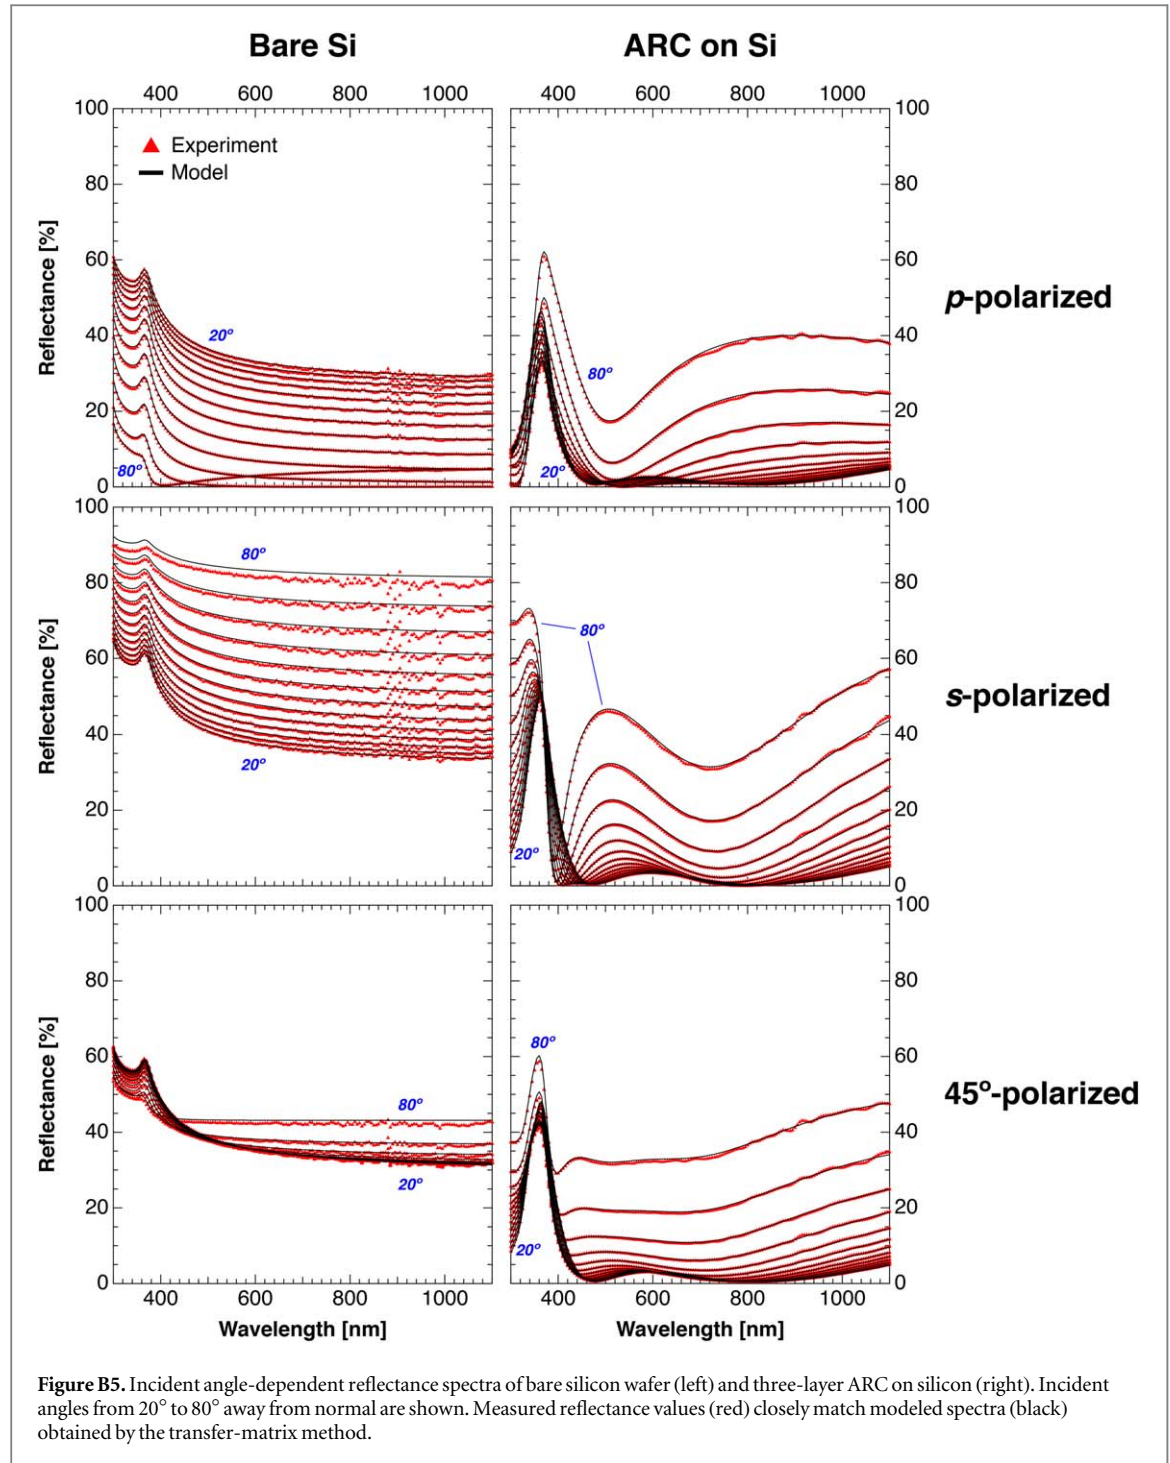
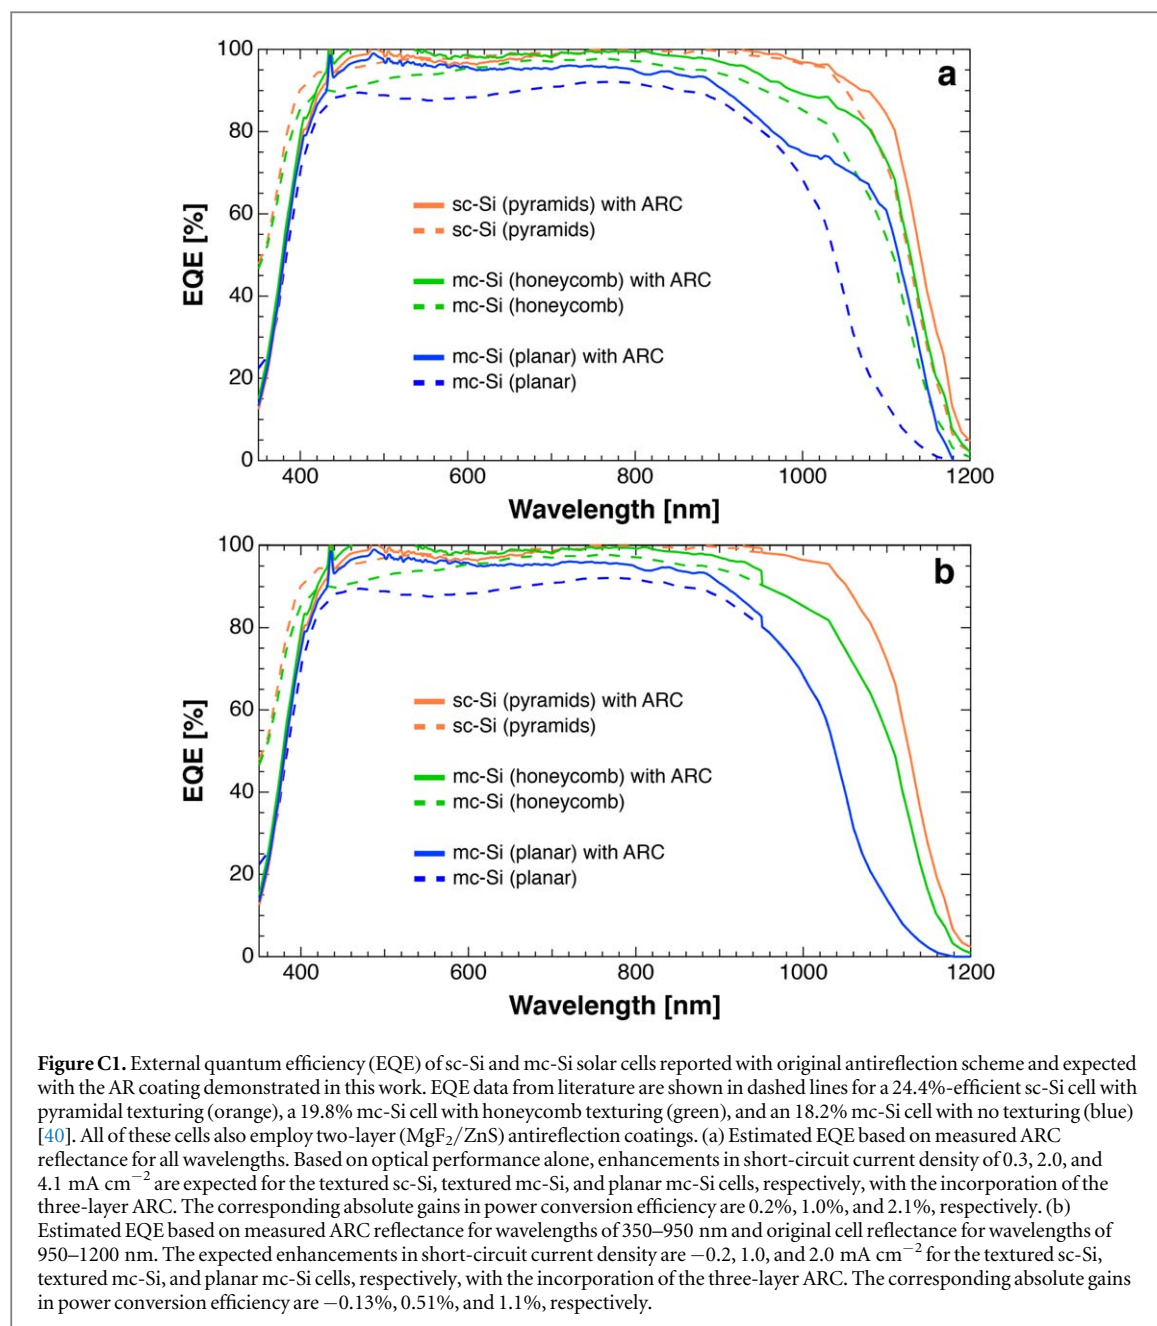


Figure B4. High-resolution broadband omnidirectional reflection from our 3-layer coating, measured using a Cary 500i spectrophotometer with a variable-angle specular reflectance accessory (average over s and p polarizations). The limits on the wavelength and incident angle (shown in units of degrees) ranges are due to instrument constraints.



Appendix C. Calculation of expected c-Si PV performance with optimized ARC

Theoretical solar cell performance with the optimized coating (yielding reflectance R_{ARC}) is calculated based on external quantum efficiency (EQE_c) and normal-incidence reflectance (R_c) data for efficient sc-Si and mc-Si cells from literature [40]. Internal quantum efficiency (IQE) spectra are first calculated for each cell, assuming no light is transmitted: $\text{IQE}(\lambda) = \frac{\text{EQE}_c(\lambda)}{(1 - R_c(\lambda))}$. Assuming the cell IQE is unchanged with different antireflection strategies, the theoretical EQE with the optimized ARC is given by $\text{EQE}_{\text{ARC}}(\lambda) = \text{IQE}(\lambda)(1 - R_{\text{ARC}}(\lambda))$. Integration of EQE_c and EQE_{ARC} against the AM1.5G photon flux spectrum yields the expected short-circuit current density of the original cell and the cell with the optimized ARC, respectively. These values can then be used to calculate the expected improvement in power conversion efficiency. All corresponding spectra are shown in figure C1.



ORCID iDs

Carmel Rotschild  <https://orcid.org/0000-0001-6400-9301>

References

- [1] Liddell H M 1981 *Computer-Aided Techniques for the Design of Multilayer Filters* (Bristol: IOP Publishing)
- [2] Macleod H A 2010 *Thin-Film Optical Filters* (Boca Raton, FL: CRC Press)
- [3] Dobrowolski J A and Kemp R A 1990 Refinement of optical multilayer systems with different optimization procedures *Appl. Opt.* **29** 2876–93
- [4] Tilsch M, Hendrix K and Verly P 2006 Optical interference coatings design contest 2004 *Appl. Opt.* **45** 1544–54
- [5] Tilsch M and Hendrix K 2008 Optical interference coatings design contest 2007: triple bandpass filter and nonpolarizing beam splitter *Appl. Opt.* **47** C55–69
- [6] Hendrix K and Oliver J 2011 Optical interference coatings design contest 2010: solar absorber and Fabry–Perot etalon *Appl. Opt.* **50** C286–300
- [7] Hendrix K, Kruschwitz J and Keck J 2014 Optical interference coatings design contest 2013: angle-independent color mirror and shortwave infrared/ midwave infrared dichroic beam splitter *Appl. Opt.* **53** A360–76
- [8] Lemarchand F 2014 Application of clustering global optimization to thin film design problems *Opt. Express* **22** 5166–76
- [9] Horst R and Tuy H 1993 *Global Optimization: Deterministic Approaches* (Berlin: Springer)

- [10] Kan A H G R and Timmer G T 1987 Stochastic global optimization methods: I. Clustering methods *Math. Program.* **39** 27–56
- [11] Smith W 2007 *Modern Optical Engineering* (New York: McGraw-Hill Education)
- [12] Martin S, Rivory J and Schoenauer M 1995 Synthesis of optical multilayer systems using genetic algorithms *Appl. Opt.* **34** 2247–54
- [13] Ghebrebrhan M, Bermel P, Avniel Y, Joannopoulos J D and Johnson S G 2009 Global optimization of silicon photovoltaic cell front coatings *Opt. Express* **17** 7505–18
- [14] Berz M and Makino K 2009 Rigorous global search using Taylor models *Proc. 2009 Conf. on Symbolic Numeric Computation* (New York: ACM) pp 11–20
- [15] Makino K and Berz M 2011 Optimal correction and design parameter search by modern methods of rigorous global optimization *Nucl. Instrum. Methods Phys. Res. A* **645** 332–7
- [16] Macdonald D H, Cuevas A, Kerr M J, Samundsett C, Ruby D, Winderbaum S and Leo A 2004 Texturing industrial multicrystalline silicon solar cells *Sol. Energy* **76** 277–83
- [17] Vazsonyi E, Clercq K D, Einhaus R, Kerschaver E V, Said K, Poortmans J, Szlufcik J and Nijs J 1999 Improved anisotropic etching process for industrial texturing of silicon solar cells *Sol. Energy Mater. Sol. Cells* **57** 179–88
- [18] Aberle A G 2001 Overview on SiN surface passivation of crystalline silicon solar cells *Sol. Energy Mater. Sol. Cells* **65** 239–48
- [19] Sopori B L and Pryor R A 1983 Design of antireflection coatings for textured silicon solar cells *Sol. Cells* **8** 249–61
- [20] Soppe W, Rieffe H and Weeber A 2005 Bulk and surface passivation of silicon solar cells accomplished by silicon nitride deposited on industrial scale by microwave PECVD *Prog. Photovolt., Res. Appl.* **13** 551–69
- [21] Southwell W H 1983 Gradient-index antireflection coatings *Opt. Lett.* **8** 584–6
- [22] Xi J Q, Schubert M F, Kim J K, Schubert E F, Chen M, Lin S Y, Liu W and Smart J A 2007 Optical thin-film materials with low refractive index for broadband elimination of Fresnel reflection *Nat. Photon.* **1** 176–9
- [23] Wang K X, Yu Z, Liu V, Cui Y and Fan S 2012 Absorption enhancement in ultrathin crystalline silicon solar cells with antireflection and light-trapping nanocone gratings *Nano Lett.* **12** 1616–9
- [24] Zhang R Y, Shao B, Dong J R, Zhang J C and Yang H 2011 Absorption enhancement analysis of crystalline Si thin film solar cells based on broadband antireflection nanocone grating *J. Appl. Phys.* **110** 1131051–8
- [25] Pathi P, Peer A and Biswas R 2017 Nano-phonic structures for light trapping in ultra-thin crystalline silicon solar cells *Nanomaterials* **7** 1–16
- [26] Huang Y et al 2007 Improved broadband and quasi-omnidirectional anti-reflection properties with biomimetic silicon nanostructures *Nat. Nanotechnol.* **2** 770–4
- [27] Jean J, Brown P R, Jaffe R L, Buonassisi T and Bulovic V 2015 Pathways for solar photovoltaics *Energy Environ. Sci.* **8** 1200–19
- [28] Azunre P 2014 A parallel branch-and-bound algorithm for thin-film optical systems, with application to realizing a broadband omnidirectional antireflection coating for silicon solar cells *PhD Thesis* Massachusetts Institute of Technology, Cambridge
- [29] Moore R E 1966 *Interval Analysis* (Philadelphia, PA: SIAM)
- [30] Rump S M 1999 INTLAB—INTERVAL LABORATORY *Developments in Reliable Computing* (Dordrecht: Kluwer)
- [31] Hofschuster W and Kramer W 2001 *C-XSC 2.0: A C++ Library for Extended Scientific Computing* (Berlin: Springer)
- [32] Hansen E and Walster G W 2004 *Global Optimization Using Interval Analysis* (New York: Marcel Dekker)
- [33] Berz M and Makino K 2011 *COSY INFINITY 9.1 Programmer's Manual* (Berlin: Springer)
- [34] Oskooi A, Mutapcic A, Noda S, Joannopoulos J D, Boyd S P and Johnson S G 2012 Robust optimization of adiabatic tapers for coupling to slow-light photonic-crystal waveguides *Opt. Express* **20** 21558–75
- [35] Minot M J 1976 Single-layer, gradient refractive index antireflection films effective from 0.35 to 2.5 micrometers *J. Opt. Soc. Am.* **6** 515–9
- [36] Kuo M, Poxson D J, Kim Y S, Mont F W, Kim J K, Schubert E F and Lin S 2008 Realization of a near-perfect antireflection coating for silicon solar energy utilization *Opt. Express* **33** 2527–9
- [37] Spinelli P, Verschuuren M A and Polman A 2012 Broadband omnidirectional antireflection coating based on subwavelength surface Mie resonators *Nat. Commun.* **3** 1–5
- [38] Ineichen P 2008 A broadband simplified version of the Solis clear sky model *Sol. Energy* **82** 758–62
- [39] Mueller R et al 2004 Rethinking satellite-based solar irradiance modeling: the SOLIS clear-sky module *Remote Sens. Environ.* **91** 160–74
- [40] Zhao J, Wang A, Green M A and Ferrazza F 1998 19.8% efficient 'honeycomb' textured multicrystalline and 24.4% monocrystalline silicon solar cells *Appl. Phys. Lett.* **73** 1991–3
- [41] Lin C and Snyder L 2008 *Principles of Parallel Programming* (Reading, MA: Addison-Wesley)