

Distributed non-equilibrium Green's function algorithms for the simulation of nanoelectronic devices with scattering

Stephen Cauley, Mathieu Luisier, Venkataramanan Balakrishnan, Gerhard Klimeck, and Cheng-Kok Koh

Citation: *J. Appl. Phys.* **110**, 043713 (2011); doi: 10.1063/1.3624612

View online: <http://dx.doi.org/10.1063/1.3624612>

View Table of Contents: <http://jap.aip.org/resource/1/JAPIAU/v110/i4>

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

Related Articles

Information processing with a single multifunctional nanofluidic diode
Appl. Phys. Lett. **101**, 133108 (2012)

Subthreshold swings below 60mV/dec in three-terminal nanojunctions at room temperature
Appl. Phys. Lett. **101**, 133504 (2012)

Electrical stabilities and carrier transport mechanisms of flexible organic bistable devices based on CdSe-InP core-shell nanoparticle/polystyrene nanocomposites
APL: Org. Electron. Photonics **5**, 203 (2012)

Electrical stabilities and carrier transport mechanisms of flexible organic bistable devices based on CdSe-InP core-shell nanoparticle/polystyrene nanocomposites
Appl. Phys. Lett. **101**, 103305 (2012)

Conductance fluctuations as a function of sliding motion in bilayer graphene nanoribbon junction: A first-principles investigation
Appl. Phys. Lett. **101**, 083101 (2012)

Additional information on *J. Appl. Phys.*

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

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

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

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

ADVERTISEMENT



AIP Advances

Special Topic Section:
PHYSICS OF CANCER

Why cancer? Why physics? [View Articles Now](#)

Distributed non-equilibrium Green's function algorithms for the simulation of nanoelectronic devices with scattering

Stephen Cauley,¹ Mathieu Luisier,² Venkataramanan Balakrishnan,² Gerhard Klimeck,² and Cheng-Kok Koh²

¹*Athinoula A. Martinos Center for Biomedical Imaging, Department of Radiology, Massachusetts General Hospital, Charlestown, Massachusetts 02129, USA*

²*School of Electrical and Computer Engineering Purdue University, West Lafayette, Indiana 47907-2035, USA*

(Received 25 March 2011; accepted 2 July 2011; published online 24 August 2011)

Through the non-equilibrium Green's function (NEGF) formalism, quantum-scale device simulation can be performed with the inclusion of electron-phonon scattering. However, the simulation of realistically sized devices under the NEGF formalism typically requires prohibitive amounts of memory and computation time. Two of the most demanding computational problems for NEGF simulation involve mathematical operations with structured matrices called semiseparable matrices. In this work, we present parallel approaches for these computational problems which allow for efficient distribution of both memory and computation based upon the underlying device structure. This is critical when simulating realistically sized devices due to the aforementioned computational burdens. First, we consider determining a distributed compact representation for the retarded Green's function matrix G^R . This compact representation is exact and allows for any entry in the matrix to be generated through the inherent semiseparable structure. The second parallel operation allows for the computation of electron density and current characteristics for the device. Specifically, matrix products between the distributed representation for the semiseparable matrix G^R and the self-energy scattering terms in $\Sigma^<$ produce the less-than Green's function $G^<$. As an illustration of the computational efficiency of our approach, we stably generate the mobility for nanowires with cross-sectional sizes as large as 4.5 nm, assuming an atomistic model with scattering. © 2011 American Institute of Physics. [doi:10.1063/1.3624612]

I. INTRODUCTION

In the absence of electron-phonon scattering, the problem of computing density of states and transmission through the non-equilibrium Green's function (NEGF) formalism reduces to a mathematical problem of finding select entries from the inverse of a typically large and sparse block tridiagonal matrix. The block tridiagonal structure examined is a direct result of the two terminal device model. The presence of a third contact would cause a deviation from this sparsity pattern, requiring corrective mathematical operations. The three contact problem is beyond the scope of this paper, but preliminary work has convinced us that techniques such as the Sherman-Morrison-Woodbury formula can be used to extend our approach for these devices. Although there has been research into numerically stable and computationally efficient serial computing algorithms,¹ when analyzing certain device geometries this type of method will result in prohibitive amounts of computation and memory consumption. In Ref. 2, a parallel divide-and-conquer algorithm (PDIV) was shown to be effective for NEGF based simulations. Two applications were presented, the atomistic level simulation of silicon nanowires and the two-dimensional simulation of nanotransistors. Alternative, serial computing, NEGF based approaches such as Ref. 3 rely on specific problem structure (currently not capable of addressing atomistic models), with computational limitations that again restrict the device size. In addition, the ability to compute information needed to determine current characteristics for devices has not been demonstrated with methods such as Ref. 3.

The prohibitive computational properties associated with NEGF based simulation has prompted a transition to wave function based methods, such as those presented in Ref. 4. Given an assumed basis structure, the problem translates from calculating select entries from the inverse of a matrix to solving large sparse systems of linear equations. This is an attractive alternative because solving large sparse systems of equations is one of the most well studied problems in applied mathematics and physics. In addition, popular algorithms such as UMFPACK⁵ and SuperLU⁶ have been constructed to algebraically (based solely on the matrix) exploit problem specific structure in an attempt to minimize the amount of computation. The performance of these algorithms for the wave function based analysis of several silicon nanowires has been examined in Ref. 7. However, for many devices of interest wave function methods are currently unable to address more sophisticated analyses that involve electron-phonon scattering. Thus, there remains a strong need to further develop NEGF based algorithms for the simulation of realistically sized devices considering these general modeling techniques. It is important to note that in this work we limit our scope to a block diagonal modeling of the scattering self-energy terms. Our method was not formulated to capture full-scattering self-energy functions.

When incorporating the effects of scattering into NEGF based simulation, it becomes necessary to determine the entire inverse of the coefficient matrix associated with the device. This substantially increases both the computational and memory requirements. There are a number of theoretical results describing the structure of the inverses of block tridiagonal

and block-banded matrices. Representations for the inverses of tridiagonal, banded, and block tridiagonal matrices can be found in Refs. 8–13. It has been shown that the inverse of a tridiagonal matrix can be compactly represented by two sequences $\{u_i\}$ and $\{v_i\}$.^{14–17} This result was extended to the cases of block tridiagonal and banded matrices in Refs. 18–20, where the $\{u_i\}$ and $\{v_i\}$ sequences generalized to matrices $\{U_i\}$ and $\{V_i\}$. Matrices that can be represented in this fashion are more generally known as semiseparable matrices.^{20,21} Typically, the computation of parameters $\{u_i\}$ and $\{v_i\}$ suffers from numerical instability even for modest-sized problems.²² It is well understood that for matrices arising in many physical applications the $\{u_i\}$ and $\{v_i\}$ sequences grow exponentially^{17,23} with the index i . One approach that has been successful in ameliorating these problems, for the tridiagonal case, is the generator approach shown in Ref. 24. Here, ratios for sequential elements of the $\{u_i\}$ and $\{v_i\}$ sequences are used as the generators for the inverse of a tridiagonal matrix. Such an approach is numerically stable for matrices of very large sizes. The extension of this generator approach to the general block-tridiagonal matrices was discussed by the same authors in Ref. 13. The authors used the block factorization of the original block-tridiagonal matrix to construct a block Cholesky decomposition of its inverse.

A generator based approach for inversion, typically referred to as the Recursive Green’s Function (RGF) algorithm, was introduced in Ref. 1 for NEGF based simulation. It is important to note that the method of Ref. 1 requires $3 \times$ less memory to compute only the density of states and transmission (in the absence of scattering), when compared to the complete generator representation. In this work, we extend the approach of Ref. 2 to consider the computation of a distributed generator representation for the inverse of the block tridiagonal coefficient matrix. We then demonstrate how the distributed generator representation allows for the efficient computation of the electron density and current characteristics of the device. Our parallel algorithms facilitate the simulation of realistically sized devices by utilizing additional computing resources to efficiently divide both the computation time and memory requirements. As an illustration, we stably generate the mobility for 4.5 nm cross-section nanowires assuming an atomistic model with scattering.

II. INVERSES OF BLOCK TRIDIAGONAL MATRICES

A block-symmetric matrix K is block tridiagonal if it has the form

$$K = \begin{pmatrix} A_1 & -B_1 & & & & & \\ -B_1^T & A_2 & -B_2 & & & & \\ & \ddots & \ddots & \ddots & & & \\ & & & & -B_{N_y-2}^T & A_{N_y-1} & -B_{N_y-1} \\ & & & & -B_{N_y-1}^T & A_{N_y} & \end{pmatrix}, \quad (1)$$

where each $A_i, B_i \in \mathbb{R}^{N_x \times N_x}$. Thus, $K \in \mathbb{R}^{N_y N_x \times N_y N_x}$, with N_y diagonal blocks of size N_x each. We will use the notation $K = \text{tri}(A_{1:N_y}, B_{1:N_y-1})$ to represent such a block tridiagonal matrix. The NEGF based simulation of nanowires using the

*sp3d5s** atomistic tight-binding model with electron-phonon scattering has been demonstrated in Ref. 25. The block tridiagonal coefficient matrix for simulation is constructed in the following way:

$$K = (EI - H - \Sigma_R^R - \Sigma_L^R - \Sigma_S^R).$$

Here, E is the energy of interest, H is the Hamiltonian containing atomistic interactions, and $\Sigma_L^R, \Sigma_R^R,$ and Σ_S^R are the left and right boundary conditions and self energy scattering terms, respectively. In order to calculate the current characteristics for the device we must first form the retarded Greens function using the fact that $KG^R = I$. A standard numerically stable mathematical representation for the inverse of this block tridiagonal matrix is dependent on two sequences of generator matrices $\{g_i^R\}, \{g_i^{\bar{R}}\}$. Here, the terms R and \bar{R} correspond to the forward and backward propagation through the device. Specifically, we can use the diagonal blocks of the inverse $\{D_i\}$ and the generators to describe the inverse a block tridiagonal matrix K in the following manner:

$$G^R = \begin{pmatrix} D_1 & D_1 g_1^{\bar{R}} & \cdots & D_1 \prod_{k=1}^{N_y-1} g_k^{\bar{R}} \\ g_1^R D_1 & D_2 & \cdots & D_2 \prod_{k=2}^{N_y-1} g_k^{\bar{R}} \\ \vdots & \vdots & \ddots & \vdots \\ \left(\prod_{k=N_y-1}^1 g_k^R \right) D_1 & \left(\prod_{k=N_y-1}^2 g_k^{\bar{R}} \right) D_2 & \cdots & D_{N_y} \end{pmatrix}, \quad (2)$$

where the diagonal blocks of the inverse, D_i , and the generator sequences satisfy the following relationships:

$$\begin{aligned} g_1^R &= A_1^{-1} B_1, \\ g_i^R &= \left(A_i - B_{i-1}^T g_{i-1}^R \right)^{-1} B_i, \quad i = 2, \dots, N_y - 1, \\ g_{N_y-1}^{\bar{R}} &= B_{N_y-1} A_{N_y}^{-1}, \\ g_i^{\bar{R}} &= B_i \left(A_{i+1} - g_{i+1}^{\bar{R}} B_{i+1}^T \right)^{-1}, \quad i = N_y - 2, \dots, 1, \\ D_1 &= \left(A_1 - g_1^{\bar{R}} B_1^T \right)^{-1}, \\ D_{i+1} &= \left(A_{i+1} - g_{i+1}^{\bar{R}} B_{i+1}^T \right)^{-1} \left(I + B_i^T D_i g_i^R \right), \quad i = 1, \dots, N_y - 2, \\ D_{N_y} &= A_{N_y}^{-1} \left(I + B_{N_y-1}^T D_{N_y-1} g_{N_y-1}^{\bar{R}} \right). \end{aligned} \quad (3)$$

The time complexity associated with determining the parametrization of G^R by the above approach is $O(N_x^3 N_y)$, with a memory requirement of $O(N_x^2 N_y)$.

A. Alternative approach for determining the compact representation

It is important to note that if the block tridiagonal portion of G^R is known, the generator sequences g_i^R and $g_i^{\bar{R}}$ can be extracted directly, i.e., without the use of entries from K through

the generator expressions (3). Examining closely the block tridiagonal portion of G^R we find the following relations:

$$\begin{aligned} D_i \bar{g}_i^R = P_i &\Rightarrow \bar{g}_i^R = D_i^{-1} P_i, \quad i = 1, \dots, N_y - 1, \\ \bar{g}_i^R D_i = Q_i &\Rightarrow \bar{g}_i^R = Q_i D_i^{-1}, \quad i = 1, \dots, N_y - 1, \end{aligned} \quad (4)$$

where P_i denotes the $(i, i + 1)$ block entry of G^R and Q_i denotes the $(i + 1, i)$ block entry of G^R . Therefore, by being able to produce the block tridiagonal portion of G^R we have all the information that is necessary to compute the compact representation.

As was alluded to in Sec. 1, direct techniques for simulation of realistic devices often require prohibitive memory and computational requirements. To address these issues we offer a parallel divide-and-conquer approach in order to construct the compact representation for G^R , i.e., the framework allows for the parallel inversion of the coefficient matrix. Specifically, we introduce an efficient method for computing the block tridiagonal portion of G^R in order to exploit the process demonstrated in (4).

III. PARALLEL INVERSION OF BLOCK TRIDIAGONAL MATRICES

The compact representation of G^R can be computed in a distributed fashion by first creating several smaller sub-matrices ϕ_i . That is, the total number of blocks for the matrix K are divided as evenly as possible among the sub-matrices. After each individual sub-matrix inverse has been computed they can be combined in a Radix-2 fashion using the matrix inversion lemma from linear algebra. Figure 1 shows both the decomposition and the two combining levels needed to form the block tridiagonal portion of G^R , assuming K has been divided into four sub-matrices. In general, if K is separated into p sub-matrices there will be $\log p$ combining levels with a total of $p - 1$ combining operations or “steps.” The notation $\phi_{i \sim j}^{-1}$ is introduced to represent the result of any combining step, through the use of the matrix inversion lemma. For example, $\phi_{1 \sim 2}^{-1}$ is the inverse of a matrix comprised of the blocks assigned to both ϕ_1 and ϕ_2 . It is important to note that using the matrix inversion lemma repeatedly to join sub-matrix inverses will result in a prohibitive amount of memory and computation for large simulation problems. This is due to the fact that at each combining step all entries would be computed and stored. Thus, the question remains on the most efficient way to produce the block tridiagonal portion of G^R , given this general decomposition scheme for the matrix K .

In this work, we introduce a mapping scheme to transform compact representations of smaller matrix inverses into the compact representation of G^R . The algorithm is organized as follows:

- Decompose the block tridiagonal matrix K into p smaller block tridiagonal matrices;
- Assign each sub-matrix to an individual CPU;
- Independently determine the compact representations associated with each submatrix;
- Gather all information that is needed to map the sub-matrix compact representations into the compact representation for G^R ;
- Independently apply the mappings to produce a portion of the compact representation for G^R on each CPU.

The procedure described above results in a “distributed compact representation” allowing for reduced memory and computational requirements. Specifically, each CPU will eventually be responsible for the elements from both the generator sequences and diagonal blocks that correspond to the initial decomposition (e.g., if ϕ_1 is responsible for blocks 1, 2, and 3 from the matrix K , the mappings will allow for the computation of $\bar{g}_{1 \dots 3}^R, \bar{g}_{1 \dots 3}^R$, and $D_{1 \dots 3}$).

In order to derive the mapping relationships needed to produce a distributed compact representation, it is first necessary to analyze how sub-matrix inverses can be combined to form the complete inverse. Consider the decomposition of the block tridiagonal matrix K into two block tridiagonal sub-matrices and a correction term, demonstrated below.

$$K = \underbrace{\begin{pmatrix} \phi_1 & & \\ & \phi_2 & \\ & & \ddots \end{pmatrix}}_{\bar{K}} + XY,$$

$$\phi_1 = \text{tri}(A_{1:k}, B_{1:k-1}), \quad \phi_2 = \text{tri}(A_{k+1:N_y}, B_{k+1:N_y-1}), \text{ and}$$

$$X = \begin{pmatrix} 0 & \dots & -B_k^T & 0 & \dots & 0 \\ 0 & \dots & 0 & -B_k & \dots & 0 \end{pmatrix}^T, \quad Y = \begin{pmatrix} 0 & \dots & 0 & I & \dots & 0 \\ 0 & \dots & I & 0 & \dots & 0 \end{pmatrix}.$$

Thus, the original block tridiagonal matrix can be decomposed into the sum of a block diagonal matrix (with its two diagonal blocks themselves being block tridiagonal) and a correction term parametrized by the $N_x \times N_x$ matrix B_k , which we will refer to as the “bridge matrix.” Using the matrix inversion lemma, we have




FIG. 1. Decomposition of block tridiagonal matrix K into four sub-matrices, where the shaded blocks correspond to the bridge matrices. The two combining levels follow the individual sub-matrix inversions, where $\phi_{i \sim j}^{-1}$ represents the inverse of sub-matrices ϕ_i through ϕ_j from the matrix K . Matrix mappings will be used to capture the combining effects and allow for the direct computation of the block tridiagonal portion of G^R .

$$G^R = (\tilde{K} + XY)^{-1} = \tilde{K}^{-1} - (\tilde{K}^{-1}X)(I + Y\tilde{K}^{-1}X)^{-1}(Y\tilde{K}^{-1}),$$

where

$$\begin{aligned} \tilde{K}^{-1}X &= \begin{pmatrix} -\phi_1^{-1}(:,k)B_k & 0 \\ 0 & -\phi_2^{-1}(:,1)B_k^T \end{pmatrix}, \\ (I + Y\tilde{K}^{-1}X)^{-1} &= \begin{pmatrix} I & -\phi_2^{-1}(1,1)B_k^T \\ -\phi_1^{-1}(k,k)B_k & I \end{pmatrix}^{-1}, \\ Y\tilde{K}^{-1} &= \begin{pmatrix} 0 & \phi_2^{-1}(:,1)^T \\ \phi_1^{-1}(:,k)^T & 0 \end{pmatrix}, \end{aligned} \quad (5)$$

and $\phi_1^{-1}(:,k)$ and $\phi_2^{-1}(:,1)$ denote, respectively, the last and first block columns of ϕ_1^{-1} and ϕ_2^{-1} .

This shows that the entries of \tilde{K}^{-1} are modified using the entries from the first rows and last columns of ϕ_1^{-1} and ϕ_2^{-1} , as well as the bridge matrix B_k . Specifically, since ϕ_1 is before or “above” the bridge point we only need the last column of its inverse to reconstruct G^R . Similarly, since ϕ_2 is after or “below” the bridge point we only need the first column of its inverse. These observations were noted in Ref. 2, where the authors demonstrated a parallel divide-and-conquer approach to determine the diagonal entries for the inverse of block tridiagonal matrices. We begin by generalizing the method from Ref. 2 in order to compute all information necessary to determine the distributed compact representation of G^R (3). That is, we would like to create a combining methodology for sub-matrix inverses with two major goals in mind. First, it must allow for the calculation of all information that would be required to repeatedly join sub-matrix inverses, in order to mimic the combining process shown in Fig. 1. Second, at the final stage of the combining process it must facilitate the computation of the block tridiagonal portion for the combined inverses. Details pertaining to the parallel computation of G^R are provided in Appendix A. The time complexity of the algorithm presented is $O(N_x^3 N_y / p + N_x^3 \log p)$, with memory consumption $O(N_x^2 N_y / p + N_x^2)$. The distribution of the compact representation is at the foundation of an efficient parallel method for calculating the less-than Green’s function $G^<$ and greater-than Green’s function $G^>$.

IV. PARALLEL COMPUTATION OF THE LESS-THAN GREEN’S FUNCTION

The parallel inversion algorithm described in Sec. III not only has advantages in computational and memory efficiency but also facilitates the formulation of a fast, and highly scalable, parallel matrix multiplication algorithm. This plays an important role during the simulation process due to the fact that computation of the less-than Green’s function requires matrix products with the retarded Green’s function matrix

$$G^< = G^R \sum^< G^{R*}. \quad (6)$$

$\Sigma^<$, which we will refer to as the less-than scattering matrix, is typically assumed to be a block diagonal matrix. We will demonstrate how the distributed compact representation of the semiseparable matrix G^R presented in Sec. III can be used to calculate the necessary information from $G^<$. Specifically, the electron density for the device will be calculated

through the diagonal entries of $G^<$ and the current characteristics through the first off-diagonal blocks of $G^<$.

A. Mathematical description

Recall that our initial state for this procedure would assume that portions of the block tridiagonal (corresponding to the size and location of the sub-matrices) of G^R have been calculated and stored. It is important to note that there are many generator representations for G^R and we would like to select a representation that will facilitate efficient calculation of $G^<$. For the mathematical operation shown in Eq. (6), our starting point will be describing the k^{th} block row of G^R in terms of D_i , $(g_i^R)^T$, and g_i^R , the diagonal blocks and two generator sequences, respectively. The following expressions are used to determine the generators from the block tridiagonal of the semiseparable matrix:

$$\begin{aligned} P_i &= g_i^R D_{i+1} \Rightarrow g_i^R = P_i (D_{i+1})^{-1}, \\ Q_i &= (g_i^R)^T D_i \Rightarrow (g_i^R)^T = Q_i (D_i)^{-1}. \end{aligned} \quad (7)$$

D_i , P_i , and Q_i are the diagonal, upper diagonal, and lower diagonal blocks, respectively. Thus, the generators $(g_i^R)^T$ and g_i^R are used to describe the k^{th} block row of G^R semiseparable matrix in the following way:

$$G^R(k,:) = \left[\prod_{i=k-1}^1 (g_i^R)^T D_1 \cdots (g_{k-1}^R)^T D_{k-1} D_k g_k^R D_{k+1} \cdots \prod_{i=k}^{N_y-1} g_i^R D_{N_y} \right].$$

This generator representation for G^R along with the block diagonal structure of $\Sigma^<$ allows for us to express each diagonal block of $G^<$ in terms of recursive sequences. Both the forward recursive sequence $g_i^<$ and backward recursive sequence $g_i^>$ are dependent on a common sequence of injections terms J_i (note: the arrow orientation for $g^<$ matches that of g^R). The relationships between the sequences and the diagonal blocks of $G^<$ are shown below.

$$\begin{aligned} J_i &= D_i \Sigma_i^< D_i^*, \quad i = 1, 2, \dots, N_y; \\ g_1^< &= J_1; \\ g_i^< &= J_i + (g_{i-1}^R)^T g_{i-1}^< (g_{i-1}^R)^C, \quad i = 2, \dots, N_y; \\ g_{N_y}^< &= J_{N_y}; \\ g_i^< &= J_i + g_i^R g_{i+1}^< (g_i^R)^*, \quad i = N_y - 1, \dots, 1; \\ G^<(1,1) &= g_1^<; \\ G^<(i,i) &= g_i^< + g_i^< - J_i, \quad i = 2, \dots, N_y. \end{aligned} \quad (8)$$

Similar serial recursions have been shown in Refs. 1 and 26. Our strategy in this work is to exploit the distributed compact representation of G^R in order to produce these sequences efficiently. That is, we will divide the computation needed to calculate the three sequences J_i , $g_i^<$, and $g_i^>$ into several sub-problems that can be efficiently separated across many processors.

As motivation, if we assume that there are $N_y = 2N$ blocks, we can define two sub-problems by separating $\Sigma^< = \Sigma^{<:1} + \Sigma^{<:2}$, where $\Sigma^{<:1}$ contains blocks 1 through N and $\Sigma^{<:2}$ contains blocks $N + 1$ through $2N$. Then, we can write

$$\begin{aligned} G^< &= G^R \Sigma^< G^{R*} = G^R (\Sigma^{<:1} + \Sigma^{<:2}) G^{R*} \\ &= G^R \Sigma^{<:1} G^{R*} + G^R \Sigma^{<:2} G^{R*} = G^{<:1} + G^{<:2}. \end{aligned}$$

For this example, we have assumed an equal separation of the diagonal blocks for the scattering matrix $\Sigma^<$, i.e., $\Sigma_i^{<:1} = 0, \forall i > N$. Thus, for the first sub-problem we have

$$\begin{aligned} J_{i:1} &= D_i \Sigma_i^{<:1} D_i^*, \quad i = 1, 2, \dots, N \\ J_{i:1} &= 0, \quad i = N + 1, 2, \dots, 2N \\ \vec{g}_{1:1}^< &= J_{1:1} \\ \vec{g}_{i:1}^< &= J_{i:1} + \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:1}^< \left(\vec{g}_{i-1}^R \right)^C, \quad i = 2, \dots, N \\ \vec{g}_{i:1}^< &= \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:1}^< \left(\vec{g}_{i-1}^R \right)^C, \quad i = N + 1, \dots, 2N \\ g_{i:1}^< &= 0, \quad i = 2N, \dots, N + 1 \\ g_{N:1}^< &= J_{N:1} \\ g_{i:1}^< &= J_{i:1} + g_i^R g_{i+1:1}^< \left(g_i^R \right)^*, \quad i = N - 1, \dots, 1. \end{aligned}$$

We then see the following relationships for the second sub-problem:

$$\begin{aligned} J_{i:2} &= 0, \quad i = 1, 2, \dots, N \\ J_{i:2} &= D_i \Sigma_i^{<:2} D_i^*, \quad i = N + 1, 2, \dots, 2N \\ \vec{g}_{i:2}^< &= 0, \quad i = 1, \dots, N \\ \vec{g}_{N+1:2}^< &= J_{N+1:2} \\ \vec{g}_{i:2}^< &= J_{i:2} + \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:2}^< \left(\vec{g}_{i-1}^R \right)^C, \quad i = N + 2, \dots, 2N \\ g_{2N:2}^< &= J_{2N:2} \\ g_{i:2}^< &= J_{i:2} + g_i^R g_{i+1:2}^< \left(g_i^R \right)^*, \quad i = 2N - 1, \dots, N + 1 \\ g_{i:2}^< &= g_i^R g_{i+1:2}^< \left(g_i^R \right)^*, \quad i = N, \dots, 1. \end{aligned}$$

The recursions for the case of two sub-problems are illustrated in Fig. 2 with $N_y = 2N = 6$. Here, the shaded terms in Figs. 2(a) and 2(b) represent the terms for each sub-problem that will be computed on CPU 1. In summary, the separation of the diagonal blocks, the generator sequences, and the scattering matrix evenly across two computers will result in the following order of operations:

	CPU1	CPU2
Stage I	$J_{i:1} = D_i \Sigma_i^{<:1} D_i^*,$ $i = 1, 2, \dots, N$	$J_{i:2} = D_i \Sigma_i^{<:2} D_i^*,$ $i = N + 1, 2, \dots, 2N$
Stage II	$\vec{g}_{1:1}^< = J_{1:1}$ $\vec{g}_{i:1}^< = J_{i:1} + \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:1}^< \left(\vec{g}_{i-1}^R \right)^C,$ $i = 2, \dots, N$	$g_{2N:2}^< = J_{2N:2}$ $g_{i:2}^< = J_{i:2} + g_i^R g_{i+1:2}^< \left(g_i^R \right)^*,$ $i = 2N - 1, \dots, N + 1$
Stage III	$g_{i:2}^< = g_i^R g_{i+1:2}^< \left(g_i^R \right)^*,$ $i = N, \dots, 1$	$\vec{g}_{i:1}^< = \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:1}^< \left(\vec{g}_{i-1}^R \right)^C,$ $i = N + 1, \dots, 2N$
Stage IV	$g_{N:1}^< = J_{N:1}$ $g_{i:1}^< = J_{i:1} + g_i^R g_{i+1:1}^< \left(g_i^R \right)^*,$ $i = N - 1, \dots, 1$	$g_{N+1:2}^< = J_{N+1:2}$ $g_{i:2}^< = J_{i:2} + \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:2}^< \left(\vec{g}_{i-1}^R \right)^C,$ $i = N + 2, \dots, 2N$

If we consider multiplication to be of order N_x^3 , then in the asymptotic sense a single processor $G^<$ calculation would require $3 \times (2N_y)N_x^3 = 6N_yN_x^3$ operations. In this case we have four stages, each with $2 \times \frac{N_y}{2}N_x^3 = N_yN_x^3$ operations. Therefore, using two processors we have reduced the number of multiplications to $4N_yN_x^3$. We now generalize this process

and demonstrate further speed-up as both the number of processors and the length of the device increase.

B. Parallel implementation

In order to simplify the presentation of the method, we will assume that the number of sub-problems p evenly divides the

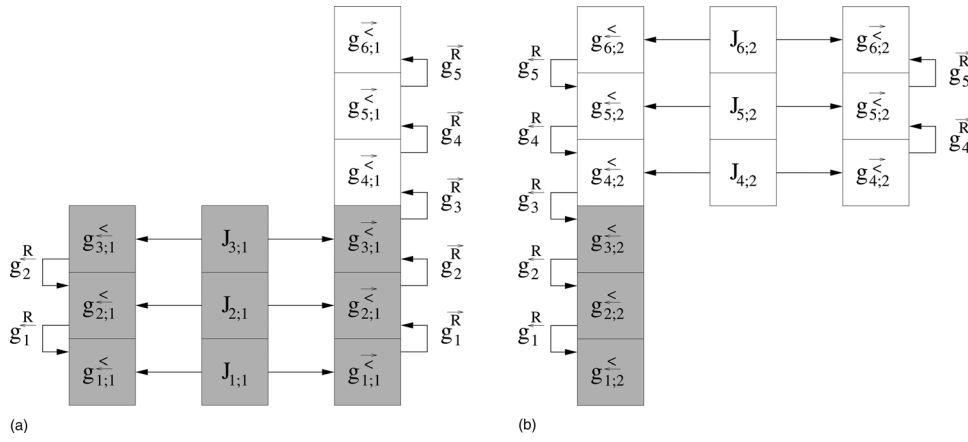


FIG. 2. Distribution of $G^<$ computation into two sub-problems across two CPUs. (a) The non-zero matrices that need to be computed for sub-problem 1. The terms that are computed on CPU 1 are shaded. (b) The non-zero matrices that need to be computed for sub-problem 2. The terms that are computed on CPU 1 are shaded.

total number of blocks $N_y = pN$. In general this assumption is not required. We will separate the $G^<$ operation evenly into p sub-problems by dividing $\Sigma^< = \Sigma^<:1 + \Sigma^<:2 + \dots + \Sigma^<:p$, where

$\Sigma^<:k$ contains the k^{th} portion of the less-than scattering matrix. Given that $\Sigma_i^<:k = 0, \forall i > kN$, and $i < (k-1)N + 1$, the k^{th} sub-problem will be solved through the following recursions:

$$\begin{aligned}
 J_{i:k} &= 0, & i &= 1, \dots, (k-1)N \\
 J_{i:k} &= D_i \Sigma_i^<:k D_i^*, & i &= (k-1)N + 1, \dots, kN \\
 J_{i:k} &= 0, & i &= kN + 1, \dots, pN \\
 \vec{g}_{i:k} &= 0, & i &= 1, \dots, (k-1)N \\
 \vec{g}_{i:k} &= J_{i:k}, & i &= (k-1)N + 1 \\
 \vec{g}_{i:k} &= J_{i:k} + \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:k}^< \left(\vec{g}_{i-1}^R \right)^C, & i &= (k-1)N + 2, \dots, kN \\
 \vec{g}_{i:k} &= \left(\vec{g}_{i-1}^R \right)^T \vec{g}_{i-1:k}^< \left(\vec{g}_{i-1}^R \right)^C, & i &= kN + 1, \dots, pN \\
 \overset{<}{g}_{i:k} &= 0, & i &= pN, \dots, kN + 1 \\
 \overset{<}{g}_{i:k} &= J_{i:k}, & i &= kN \\
 \overset{<}{g}_{i:k} &= J_{i:k} + g_i^R \overset{<}{g}_{i+1:k}^< \left(g_i^R \right)^*, & i &= kN - 1, \dots, (k-1)N + 1 \\
 \overset{<}{g}_{i:k} &= g_i^R \overset{<}{g}_{i+1:k}^< \left(g_i^R \right)^*, & i &= (k-1)N, \dots, 1.
 \end{aligned}$$

By splitting the scattering matrix we have divided the injection sequence J evenly and created two new sub-problem propagating sequences $\vec{g}_{i:k}^<$ and $\overset{<}{g}_{i:k}$. However, the sub-problem propagating sequences have the special property that they do not start (are zero) until they reach the indices governed by the sub-problem. This is due to the fact that there are no injection terms for the given sub-problem until these points are reached. In addition, once outside the range of the sub-problem we do not have any additional injection terms in the recursions. Thus, we have a standard two-sided autoregressive expression where the terms are simply propagated by multiplication with generator matrices.

Recursions for the case of four sub-problems are illustrated in Fig. 3 for an example with $N_y = 4N = 12$. Here, the shaded terms in Figs. 3(a) and 3(b) represent the terms for each sub-problem that are computed on CPU 1 and CPU 4, respectively. For each sub-problem we introduce two new sequences that will be referred to as “skip matrices.” Specifically, we define for each processor k two matrices $[g^R]_k$ and $[g^R]_k^T$ that are accumulations of the generator terms stored on the processor. As can be clearly seen from Fig. 3, these skip matrices allow for several steps in the recursive process to be preformed through a single operation. Therefore, if the skip matrices are made available to each processor several terms

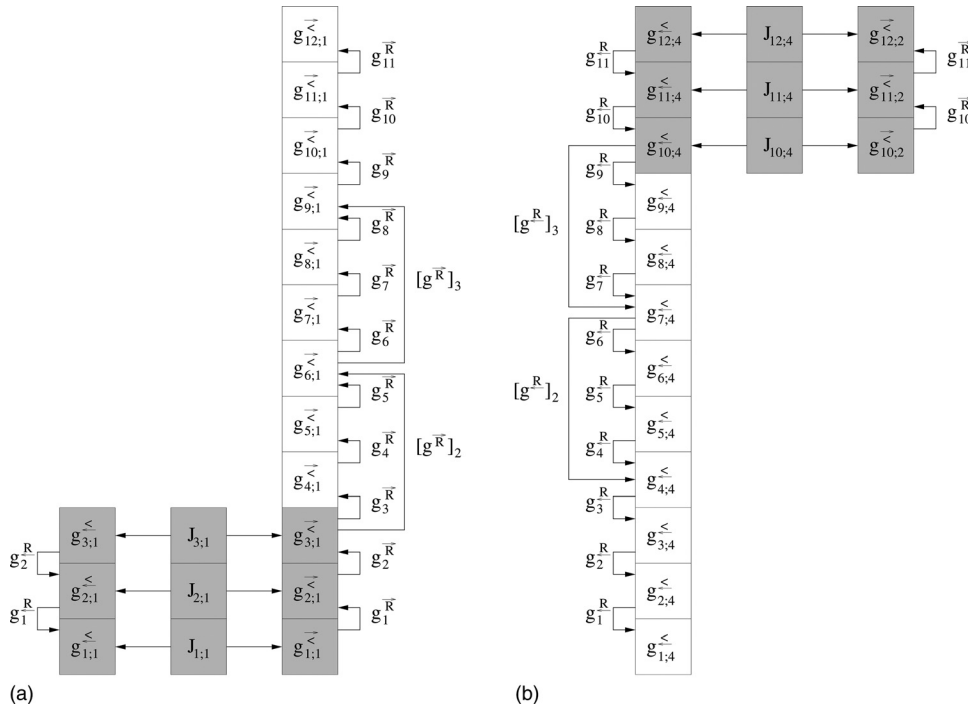


FIG. 3. Distribution of $G^<$ computation into four sub-problems across four CPUs. (a) The non-zero matrices that need to be computed for sub-problem 1. The terms that are computed on CPU 1 are shaded. (b) The non-zero matrices that need to be computed for sub-problem 4. The terms that are computed on CPU 4 are shaded.

for each sub-problem can be found concurrently. In Fig. 3(b), we can see that skip matrices will allow for $g_{9,4}^<, g_{6,4}^<$, and $g_{3,4}^<$ to be determined currently by CPUs 3, 2, and 1, respectively.

Given that each sub-problem will produce both a forward and backward propagating sequence it should be clear that CPU k will need to perform computation for $g^<$ sequences from sub-problems $\{k+1, k+2, \dots, p\}$ and $g^>$ sequences from sub-problems $\{1, 2, \dots, k-1\}$. However, each $g^>$ sequence and $g^<$ sequence from other sub-problems may be combined before the generators on that CPU are applied. This is due to the fact that the sequences from each sub-problem will eventually be added together to form $G^< = \sum_{k=1}^p G^{<:k}$. For the example shown in Fig. 3 we see that CPU 1 will first need to form

$$g_{4,4}^< + g_{4,3}^< + g_{4,2}^< = [g^R]_2 \left([g^R]_3 g_{10,4}^< [g^R]_3^* + g_{7,3}^< \right) \times [g^R]_2^* + g_{4,2}^<$$

before the required terms for each sub-problem can be calculated. That is, after the lead term for the backward propagating sequence $(g_{4,4}^< + g_{4,3}^< + g_{4,2}^<)$ has been computed, the generators governed by CPU 1: g_3^R, g_2^R , and g_1^R , can be applied to fulfill the sub-problem solutions. We now have all the tools necessary to construct a general procedure:

1. Compute the skip products of each generator sequence g^R and $g^<$, i.e., compute $[g^R]_k = \prod_{i=(k-1)N+1}^{kN} g_i^R$ and $[g^<]_k^T = \prod_{i=kN-1}^{(k-1)N} (g_i^<)^T$, where we will define $(g_0^<)^T = g_{N_y}^R = I$.
2. Compute the initial injection and propagation terms for each sub-problem,

$$J_{i:k} = D_i \Sigma_i^{<:k} D_i^*, \quad i = (k-1)N+1, \dots, kN$$

$$g_{i:k}^< = J_{i:k}, \quad i = (k-1)N+1$$

$$g_{i:k}^< = J_{i:k} + (g_{i-1}^<)^T g_{i-1;k}^< (g_{i-1}^<)^C, \quad i = (k-1)N+2, \dots, kN$$

$$g_{i:k}^< = J_{i:k}, \quad i = kN$$

$$g_{i:k}^< = J_{i:k} + g_i^R g_{i+1;k}^< (g_i^R)^*, \quad i = kN-1, \dots, (k-1)N+1.$$

3. Transfer forward and backward skip products: $[g^R]_k^T$ and $[g^R]_k$, as well as forward and backward lead propagating matrices: $g_{kN;k}^<$ and $g_{(k-1)N+1;k}^<$, for each sub-problem k to all CPUs. This will be a total of $4pN\lambda^2$ entries.
4. CPU k will construct combined lead propagating term for $g^<$ sequences from sub-problems $\{k+1, k+2, \dots, p\}$.

$$\sum_{j=k+1}^p g_{kN+1;j}^< = g_{kN+1;k+1}^< + \sum_{j=k+2}^p \left(\prod_{l=k+1}^{j-1} [g^R]_l \right) g_{(j-1)N+1;j}^< \left(\prod_{l=j-1}^{k+1} [g^R]_l^* \right).$$

CPU k will construct combined lead propagating term for $g^<$ sequences from sub-problems $\{1, 2, \dots, k-1\}$.

$$\sum_{j=1}^{k-1} g_{(k-1)N;j}^< = g_{(k-1)N;k-1}^< + \sum_{j=1}^{k-2} \left(\prod_{l=k-1}^{j+1} [g^R]_l^T \right) g_{jN+1;j}^< \left(\prod_{l=j+1}^{k-1} [g^R]_l^C \right).$$

5. Successively multiply by the governed generator terms g^R and $g^<$ in order fulfill sub-problem solutions.

$$\sum_{j=k+1}^p g_{ij}^{\leq} = \left(\prod_{l=i}^{kN} g_l^R \right) \sum_{j=k+1}^p g_{kN+1;j}^{\leq} \left(\prod_{l=kN}^i (g_l^R)^* \right), \quad \frac{p}{2.34 + \frac{p(p-2)}{3N_y}}. \quad (10)$$

$$i = kN, \dots, (k-1)N + 1.$$

$$\sum_{j=1}^{k-1} g_{ij}^{\leq} = \left(\prod_{l=(i-1)}^{(k-1)N} (g_l^R)^T \right) \sum_{j=1}^{k-1} g_{(k-1)N;j}^{\leq} \left(\prod_{l=(k-1)N}^{(i-1)} (g_l^R)^C \right).$$

$$i = (k-1)N + 1, \dots, kN.$$

6. Each CPU will combine portions of all sub-problem solutions in order to produce corresponding portion of the diagonal blocks of G^{\leq} based upon the relationships shown in (8).

In order to analyze the computational improvement for the approach we consider the number of multiplications required for each stage. The accumulation of the skip matrices described in Step 1 will result in $2N_x^3 N_y/p$ multiplications. The individual sub-problem recursions of Step 2 results in $6N_x^3 N_y/p$ multiplications. Step 4 describes how the skip matrices may be applied in order to create the sum for each sub-problem propagating sequence, requiring $2N_x^3(p-2)$ multiplications. Finally, these two sums must be propagated through each governed generator matrices, requiring $4N_x^3 N_y/p$ multiplications. If one is interested in also computing the current characteristics for the device, they may be determined through the off-diagonal blocks of G^{\leq} ,

$$G^{\leq}(i, i+1) = g_{i+1}^{\leq} \left(g_i^R \right)^* + \left(g_i^R \right)^T g_i^{\leq}, \quad i = 1, \dots, N_y - 1. \quad (9)$$

As portions of each generator sequence and propagating sequence have been evenly distributed among the processors, the resulting computation would be $2N_x^3 N_y/p$. Therefore, the total number of multiplications for the approach is $2N_x^3(p-2) + 14N_x^3 N_y/p$ compared to $6N_x^3 N_y$ for the serial algorithm. Therefore, the speedup of our approach is

We can clearly see that if $N_y \gg p$ we will approach a speedup of $p/2.34$.

V. RESULTS

The parallel G^R and G^{\leq} algorithms have been implemented in C using MPI for inter-processor communication. All computational complexity analyses were performed on a cluster of Intel E5410 processors with 16 GB of shared memory for the 8 core machines. The OMEN simulator²⁵ was used to perform simulations for square silicon nanowires employing the *sp3d5s** atomistic tight-binding model with electron-phonon scattering. We will begin by demonstrating the computational efficiency of our approach. We will then analyze device characteristics for large cross-section nanowires that have prohibitive memory requirements for the serial RGF approach.

Each NEGF nanowire simulation requires thousands of G^R , G^{\leq} , and G^{\geq} computations. For our approach, as well as RGF, the time for each computation will remain constant given that the underlying structure of the Hamiltonian and scattering matrices does not change. In order to analyze the computational benefits of our approach we have examined several different cross-section sizes and lengths of nanowires. Specifically, silicon [100] nanowires with lengths of nearly 35 nm, 53 nm, and 70 nm are examined with cross-sections of 2, 3, and 4 nm. The corresponding block sizes for the Hamiltonian matrices will be $N_x \approx \{320, 720, 1130\}$ with number of blocks, $N_y = \{256, 384, 512\}$. Table I shows the runtime for the G^R and G^{\leq} computations (the time needed for G^{\geq} is identical to that of G^{\leq}). “RGF×” is the observed speed-up for the combined time of G^R , G^{\leq} , and G^{\geq} calculations compared to RGF. “N/A” is used for devices that are not long enough to be divided based upon the number of CPUs. We can clearly see that the efficiency of our algorithm improves as both the length and cross-section of

TABLE I. Runtime comparisons between parallel G^R and G^{\leq} algorithms and serial RGF approach. Silicon nanowires with lengths of nearly 35, 53, and 70 nm are examined with cross sections of 2, 3, and 4 nm. “RGF×” is the observed speed-up for the combined time. “N/A” is used for devices that are too short to support the number of CPUs.

p	2 nm cross section, $N_x \approx 320$				3 nm cross section, $N_x \approx 720$				4 nm cross section, $N_x \approx 1130$			
	4	8	16	32	4	8	16	32	4	8	16	32
	35nm length, $N_y = 256$											
G^R	35.8	21.5	16.2	N/A	411.1	241.4	163.6	N/A	1915.5	1099.7	726.1	N/A
G^{\leq}	22.7	12.7	10.0	N/A	253.7	133.7	81.2	N/A	1159.2	598.5	349.7	N/A
RGF×	1.5	2.7	3.4	N/A	1.4	2.6	4.1	N/A	1.4	2.6	4.2	N/A
	53 nm length, $N_y = 384$											
G^R	51.4	29.3	19.7	N/A	594.4	331.9	211.4	N/A	2756.1	1523.4	946.4	N/A
G^{\leq}	33.5	18.5	13.8	N/A	380.9	198.5	113.8	N/A	1741.6	887.8	489.1	N/A
RGF×	1.6	2.8	4.0	N/A	1.5	2.7	4.5	N/A	1.4	2.7	4.6	N/A
	70 nm length, $N_y = 512$											
G^R	67.1	38.8	23.8	17.7	779.0	422.6	255.0	181.3	3601.3	1953.7	1154.5	782.1
G^{\leq}	44.6	24.4	15.8	12.9	507.6	261.5	145.5	103.4	2324.2	1182.2	640.2	408.9
RGF×	1.6	2.8	4.5	5.7	1.5	2.8	4.8	6.8	1.4	2.7	4.9	7.4

the device increase. There are two reasons behind this trend. First, considering a fixed number of processors a longer device will devote less of the total time to sub-problem combining. In addition, as the cross-section size increases the inter-processor communication costs will be a smaller fraction of the total simulation time. We see these effects again when examining Fig. 4. Here, we verify the accuracy of the estimated $G^<$ scaling trend that was derived in Eq. (10). The speed-up over RGF for both 53 and 105 nm nanowires are compared against our theoretical estimate. It is important to note that although the speed-up estimate (10) is independent of the cross-section size, effects such as data access time, vector scaling/addition, and inter-processor communication will play a role in determining the efficiency. Thus, in both Figs. 4(a) and 4(b) we again see improved efficiency when considering the larger 3 nm cross-section nanowires. In the case of

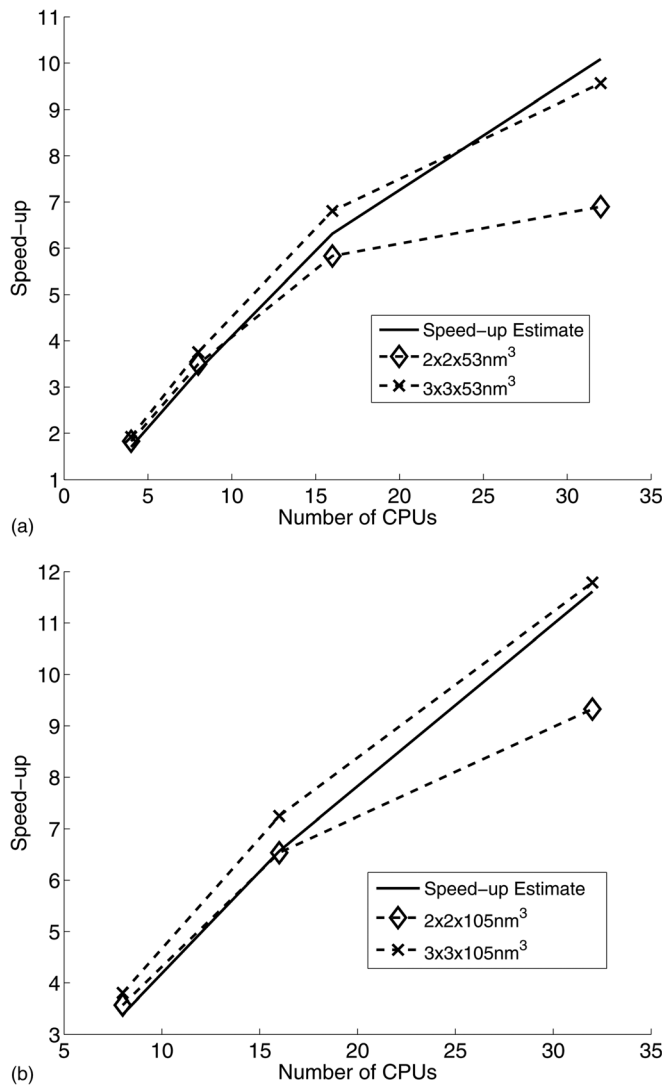


FIG. 4. Verification of theoretical estimate for $G^<$ RGF speed-up considering 53 and 105 nm length silicon nanowires. (a) $G^<$ speed-up is compared against theoretical estimate from (10) for 53 nm length silicon nanowires. The corresponding block sizes for the Hamiltonian matrices will be $N_x = \{320, 720\}$ with number of blocks $N_y = 384$. (b) $G^<$ speed-up is compared against theoretical estimate from (10) for 105 nm length silicon nanowires. The corresponding block sizes for the Hamiltonian matrices will be $N_x = \{320, 720\}$ with number of blocks $N_y = 768$.

the $3 \times 3 \times 105 \text{ nm}^3$ nanowire we achieve $11.8 \times$ speed-up when utilizing 32 CPUs.

In addition to providing computational improvements, our algorithm facilitates the simulation of larger cross-section devices. If we consider the 4 nm cross-section devices analyzed in Table I, the RGF method would require between 16 and 32 GB of memory for lengths of 35 to 70 nm. These devices would not be able to be analyzed without the use of special purpose hardware. As an illustration of the capacity for our algorithm to simulate devices previously viewed to have prohibitive memory requirements, we stably generate the mobility for 4.5 nm cross-section devices with electron-phonon scattering. In order to facilitate complete nanowire simulations we have implemented our algorithm on dual hex-core AMD Opteron 2435 (Istanbul) processors running at 2.6 GHz, 16GB of DDR2-800 memory, and a SeaStar 2+ router. As an application, the low-field phonon-limited mobility of electrons μ_{ph} is calculated in circular nanowires with diameters ranging from 1.5 to 4.5 nm and transport along the [100] crystal axis. The corresponding block sizes for the Hamiltonian matrices will be $N_x \approx \{120, 230, 350, 490, 660, 860, 1100\}$. “dR/dL” method²⁷ and the same procedure as in Ref. 28 are used to obtain μ_{ph} . The channel resistance “R” is computed as function of the nanowire length “L” and then converted into a mobility. Here, “L” is set to 35 nm ($N_y = 256$), “R” is computed in the limit of ballistic transport and in the presence of electron-phonon scattering, and the difference between these two points is considered to evaluate dR/dL. The results are shown in Fig. 5. From a numerical perspective, the computation of each Green’s function, at a given energy, was parallelized on 16 CPUs for all the device structures. As was alluded to above the simulation of these structures would not have been possible (due to memory restrictions) without the decomposition of the device through our parallel methods.




FIG. 5. Phonon-limited electron mobility μ_{ph} as function of the diameter of [100] oriented circular nanowires at room temperature. The electron density along the channel is homogeneous and set to $n = 1e20 \text{ cm}^{-3}$.

VI. CONCLUSIONS

In this work we have developed algorithms for parallel NEGF simulation with scattering. The computational benefits of our approach have been demonstrated on large cross-section silicon nanowires. We show improvements of over $11\times$ for the $G^<$ and $G^>$ computations. In addition, our approach enables simulations without the need for special purpose hardware. This can best be observed through our simulation results for 4.5 nm cross-section silicon nanowires. The algorithms developed in this work are applicable for a wide range of device geometries considering both atomistic and effective-mass models. In addition to offering significant computational improvements over the serial Recursive Green's Function algorithm, our approach facilitates simulation of realistically sized devices on typical distributed computing hardware.

APPENDIX: PARALLEL COMPUTATION OF THE RETARDED GREEN'S FUNCTION

In Appendix, we build upon the approach of Ref. 2 to consider the computation of a distributed generator representation for G^R . In Sec. 1, we introduce a mapping framework for combining sub-matrices corresponding to subsets of the device structure. This has similarities to the approach presented in Ref. 2, where in that case only the diagonal entries of G^R were needed to describe the density of states. Section 2 illustrates how a recursive process can be formulated in order to reconstruct the compact representation of G^R . This involves several key differences when compared to Ref. 2 as significantly more information is required from G^R . Finally, in Sec. 3 the parallel G^R algorithm is summarized including an analysis of the computational complexity.

1. Matrix maps

Matrix mappings are constructed in order to eventually produce the block tridiagonal portion of G^R while avoiding any unnecessary computation during the combining process. Specifically, we will show that both the boundary block entries (first block row and last block column) and the block tridiagonal entries from any combined inverse $\phi_{i\sim j}^{-1}$ must be attainable (not necessarily computed) for all combining steps. We begin by illustrating the initial stage of the combining process given four sub-matrices, where for simplicity each will be assumed to have N blocks of size N_x . First, the two sub-matrices ϕ_1 and ϕ_2 are connected through the bridge matrix B_N and together they form the larger block tridiagonal matrix $\phi_{1\sim 2}$. By examining Fig. 1, it can be seen that eventually $\phi_{1\sim 2}^{-1}$ and $\phi_{3\sim 4}^{-1}$ will be combined and we must therefore produce the boundaries for each combined inverse. From Eq. (5) the first block row and last block column of $\phi_{1\sim 2}^{-1}$ can be calculated through the use of an "adjustment" matrix,

$$J = \begin{pmatrix} I & -\phi_2^{-1}(1,1)B_N^T \\ -\phi_1^{-1}(N,N)B_N & I \end{pmatrix}^{-1},$$

as follows:

$$\begin{aligned} \phi_{1\sim 2}^{-1}(1, \cdot) &= [\phi_1^{-1}(1, \cdot) \ 0] - \left\{ \begin{array}{l} [-\phi_1^{-1}(1, N)B_N J_{12} \phi_1^{-1}(\cdot, N)^T]^T \\ [-\phi_1^{-1}(1, N)B_N J_{11} \phi_2^{-1}(1, \cdot)]^T \end{array} \right\}^T, \\ \phi_{1\sim 2}^{-1}(\cdot, 2N) &= [0 \ \phi_2^{-1}(N, \cdot)]^T - \left\{ \begin{array}{l} [-\phi_2^{-1}(N, 1)B_N^T J_{22} \phi_1^{-1}(\cdot, N)^T]^T \\ [-\phi_2^{-1}(N, 1)B_N^T J_{21} \phi_2^{-1}(1, \cdot)]^T \end{array} \right\}. \end{aligned} \quad (\text{A1})$$

In addition, the r^{th} diagonal block of $\phi_{1\sim 2}^{-1}$ can be calculated using the following relationships for $r \leq N$:

$$\begin{aligned} \phi_{1\sim 2}^{-1}(r, r) &= \phi_1^{-1}(r, r) - (-\phi_1^{-1}(r, N)B_N J_{12} \phi_1^{-1}(r, N)^T), \\ \phi_{1\sim 2}^{-1}(r+N, r+N) &= \phi_2^{-1}(r, r) - (-\phi_2^{-1}(r, 1)B_N^T J_{21} \phi_2^{-1}(1, r))^T, \end{aligned} \quad (\text{A2})$$

where the r^{th} off-diagonal block of $\phi_{1\sim 2}^{-1}$ can be calculated using the following relationships:

$$\begin{aligned} \phi_{1\sim 2}^{-1}(r, r+1) &= \phi_1^{-1}(r, r+1) - [-\phi_1^{-1}(r, N)B_N J_{12} \phi_1^{-1}(r+1, N)^T], \\ \phi_{1\sim 2}^{-1}(r+N, r+1+N) &= \phi_2^{-1}(r, r+1) - [-\phi_2^{-1}(r+1, 1)B_N^T J_{21} \phi_2^{-1}(1, r)^T], \\ r < N, \\ \phi_{1\sim 2}^{-1}(r, r+1) &= 0 - [-\phi_1^{-1}(r, N)B_N J_{11} \phi_2^{-1}(1, 1)], \\ r = N. \end{aligned} \quad (\text{A3})$$

The combination of ϕ_3 and ϕ_4 through the bridge matrix B_{3N} results in similar relationships to those seen above. Thus, in order to be able to produce both the boundary and block tridiagonal portions of each combined inverse we assign a total of twelve $N_x \times N_x$ matrix maps for each sub-matrix k . $M_{k;1-4}$ describe effects for the k^{th} portion of the boundary, $M_{k;5-8}$ describe the effects for a majority of the tridiagonal blocks, while $C_{k;1-4}$, which we will refer to as "cross" maps, can be used to produce the remainder of the tridiagonal blocks.

Initially, for each sub-matrix i the mappings $M_{k;i} = I$, $k = 1, 4$, with all remaining mapping terms set to zero. This ensures that initially the boundary of $\phi_{i\sim i}^{-1}$ matches the actual entries from the sub-matrix inverse, and the modifications to the tridiagonal portion due to combining are all set to zero. By examining the first block row, last block column, and the tridiagonal portion of the combined inverse $\phi_{1\sim 2}^{-1}$ we can see how the maps can be used to explicitly represent all of the needed information. The governing responsibilities of the individual matrix maps are detailed below.

$$\begin{aligned}
\phi_{1\sim 2}^{-1}(1, :) &= \left\{ \begin{array}{l} [M_{1;1}\phi_1^{-1}(1, :) + M_{2;1}\phi_1^{-1}(:, N)^T]^T \\ [M_{1;2}\phi_2^{-1}(1, :) + M_{2;2}\phi_2^{-1}(:, N)^T]^T \end{array} \right\}^T, \\
\phi_{1\sim 2}^{-1}(:, 2N) &= \left\{ \begin{array}{l} [M_{3;1}\phi_1^{-1}(1, :) + M_{4;1}\phi_1^{-1}(:, N)^T]^T \\ [M_{3;2}\phi_2^{-1}(1, :) + M_{4;2}\phi_2^{-1}(:, N)^T]^T \end{array} \right\}^T, \\
\phi_{1\sim 2}^{-1}(r, s) &= \phi_1^{-1}(r, s) - [\phi_1^{-1}(r, 1)M_{5;1}\phi_1^{-1}(1, s) + \phi_1^{-1}(r, 1)M_{6;1}\phi_1^{-1}(s, N)^T \\
&\quad + \phi_1^{-1}(r, N)M_{7;1}\phi_1^{-1}(1, s) + \phi_1^{-1}(r, N)M_{8;1}\phi_1^{-1}(s, N)^T], \\
\phi_{1\sim 2}^{-1}(r, s + N) &= -[\phi_1^{-1}(r, 1)C_{1;1}\phi_2^{-1}(1, s) + \phi_1^{-1}(r, 1)C_{2;1}\phi_2^{-1}(s, N)^T \\
&\quad + \phi_1^{-1}(r, N)C_{3;1}\phi_2^{-1}(1, s) + \phi_1^{-1}(r, N)C_{4;1}\phi_2^{-1}(s, N)^T], \\
\phi_{1\sim 2}^{-1}(r + N, s + N) &= \phi_2^{-1}(r, s) - [\phi_2^{-1}(r, 1)M_{5;2}\phi_2^{-1}(1, s) + \phi_2^{-1}(r, 1)M_{6;2}\phi_2^{-1}(s, N)^T \\
&\quad + \phi_2^{-1}(r, N)M_{7;2}\phi_2^{-1}(1, s) + \phi_2^{-1}(r, N)M_{8;2}\phi_2^{-1}(s, N)^T], \\
r, s &\leq N.
\end{aligned} \tag{A4}$$

It is important to note that all of the expressions (A1)–(A3) can be written into the matrix map framework of (A4). Figure 6 shows the mapping dependencies for the first block row and last block row (or column since K is symmetric). From (A1) we see that both of the block rows are distributed based upon the location of each sub-matrix with respect to the bridge point, i.e., the mapping terms associated with ϕ_1^{-1} can be used to produce the first portion of the rows while those associated with ϕ_2^{-1} can be used for the remainder. In fact, this implicit division for the mapping dependencies holds for the block tridiagonal portion of the combined inverses as well, enabling an efficient parallel implementation. Thus, from this point we can deduce that the matrix maps for the first block row (A4) must be updated in the following manner:

$$\begin{aligned}
M_{1;1} &\leftarrow M_{1;1} + (\phi_1^{-1}(1, N)B_N J_{12})M_{3;1}; \\
M_{2;1} &\leftarrow M_{2;1} + (\phi_1^{-1}(1, N)B_N J_{12})M_{4;1}; \\
M_{3;1} &\leftarrow (\phi_1^{-1}(1, N)B_N J_{11})M_{1;2}; \\
M_{4;1} &\leftarrow (\phi_1^{-1}(1, N)B_N J_{11})M_{2;2}.
\end{aligned}$$

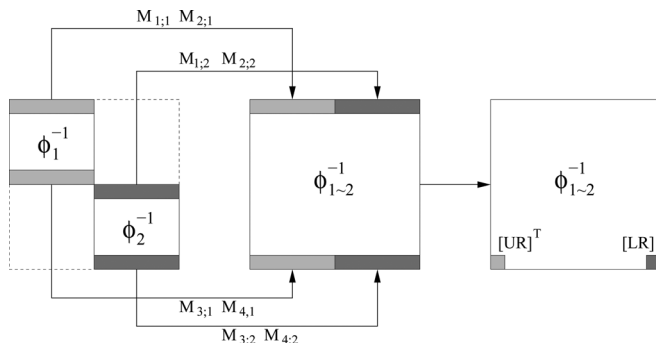


FIG. 6. Mapping dependencies when combining ϕ_1^{-1} and ϕ_2^{-1} to form $\phi_{1\sim 2}^{-1}$.

In order to understand these relationships it is important to first recall that the updates to the maps associated with sub-matrix ϕ_1 are dependent on the last block column $\phi_1^{-1}(:, N)$. Thus, we see a dependence on the previous state for the last block column $\phi_1^{-1}(:, N)$, i.e., the new state of the mapping terms $M_{1;1}$ and $M_{2;1}$ are dependent on the previous state of the mapping terms $M_{3;1}$ and $M_{4;1}$, respectively. Similarly, a dependence on $\phi_2^{-1}(1, :)$ results in the new state of the mapping terms $M_{1;2}$ and $M_{2;2}$ being functions of the previous state of the mapping terms $M_{3;2}$ and $M_{4;2}$, respectively. Finally, although some of the mapping terms remain zero after this initial combining step ($M_{2;2}$ for example), the expressions described in (A4) need to be general enough for the methodology. That is, the mapping expressions must be able to capture combining effects for multiple combining stages, regardless of the position of the sub-matrix with respect to a bridge point. For example, if we consider sub-matrix ϕ_2 for the case seen in Fig. 1, during the initial combining step it would be considered a lower problem and for the final combining step it would be considered an upper problem. Alternatively, sub-matrix ϕ_3 would be associated with exactly the opposite modifications. It is important to note that every possible modification process, for the individual mapping terms, is encompassed within this general matrix map framework.

2. Recursive combining process

In order to formalize the notion of a recursive update scheme we will continue the example from Sec. 1. By examining the final combining stage for the case of four sub-matrices, we notice that the approach described in Eqs. (A1)–(A3) can again be used to combine sub-matrix inverses $\phi_{1\sim 2}^{-1}$ and $\phi_{3\sim 4}^{-1}$, through the bridge matrix B_{2N} . The first block row and last block column of $\phi_{1\sim 4}^{-1}$ can be calculated as follows:

$$\begin{aligned}\phi_{1\sim 4}^{-1}(1, :) &= [\phi_{1\sim 2}^{-1}(1, :) \ 0] - \left\{ \begin{array}{l} [-\phi_{1\sim 2}^{-1}(1, 2N)B_{2N}J_{12}\phi_{1\sim 2}^{-1}(:, 2N)^T]^T \\ [-\phi_{1\sim 2}^{-1}(1, 2N)B_{2N}J_{11}\phi_{3\sim 4}^{-1}(1, :)]^T \end{array} \right\}^T, \\ \phi_{1\sim 4}^{-1}(:, 4N) &= [0 \ \phi_{3\sim 4}^{-1}(2N, :)]^T - \left\{ \begin{array}{l} [-\phi_{3\sim 4}^{-1}(2N, 1)B_{2N}^TJ_{22}\phi_{1\sim 2}^{-1}(:, 2N)^T]^T \\ [-\phi_{3\sim 4}^{-1}(2N, 1)B_{2N}^TJ_{21}\phi_{3\sim 4}^{-1}(1, :)]^T \end{array} \right\}^T,\end{aligned}\quad (\text{A5})$$

given the adjustment matrix:

$$J = \begin{pmatrix} I & -\phi_{3\sim 4}^{-1}(1, 1)B_{2N}^T \\ -\phi_{1\sim 2}^{-1}(2N, 2N)B_{2N} & T \end{pmatrix}^{-1}.$$

In addition, the r^{th} diagonal block of $\phi_{1\sim 4}^{-1}$ can be calculated using the following relationships:

$$\begin{aligned}\phi_{1\sim 4}^{-1}(r, r) &= \phi_{1\sim 2}^{-1}(r, r) \\ &\quad - \left(-\phi_{1\sim 2}^{-1}(r, 2N)B_{2N}J_{12}\phi_{1\sim 2}^{-1}(r, 2N)^T \right), \\ \phi_{1\sim 4}^{-1}(r + 2N, r + 2N) &= \phi_{3\sim 4}^{-1}(r, r) \\ &\quad - \left(-\phi_{3\sim 4}^{-1}(r, 1)B_{2N}^TJ_{21}\phi_{3\sim 4}^{-1}(1, r) \right)^T, \\ r &\leq 2N,\end{aligned}\quad (\text{A6})$$

where the r^{th} off-diagonal block of $\phi_{1\sim 4}^{-1}$ can be calculated using the following relationships:

$$\begin{aligned}\phi_{1\sim 4}^{-1}(r, r + 1) &= \phi_{1\sim 2}^{-1}(r, r + 1) \\ &\quad - \left[\phi_{1\sim 2}^{-1}(r, 2N)B_{2N}J_{12}\phi_{1\sim 2}^{-1}(r + 1, 2N)^T \right], \\ \phi_{1\sim 4}^{-1}(r + 2N, r + 1 + 2N) &= \phi_{3\sim 4}^{-1}(r, r + 1) \\ &\quad - \left[-\phi_{3\sim 4}^{-1}(r + 1, 1)B_{2N}^TJ_{21}\phi_{3\sim 4}^{-1}(1, r) \right]^T, \\ r &< 2N, \\ \phi_{1\sim 4}^{-1}(r, r + 1) &= 0 - \left[-\phi_{1\sim 2}^{-1}(r, 2N)B_{2N}J_{11}\phi_{3\sim 4}^{-1}(1, 1) \right], \\ r &= 2N.\end{aligned}\quad (\text{A7})$$

Again, it is important to note that each of the expressions (A5)–(A7) are implicitly divided based upon topology. For example, the first $2N$ diagonal blocks of $\phi_{1\sim 4}^{-1} = G^R$ can be separated into two groups based upon the size of the sub-matrices ϕ_1 and ϕ_2 . That is,

$$\phi_{1\sim 4}^{-1}(r, r) = \phi_{1\sim 2}^{-1}(r, r) - \left[-\phi_{1\sim 2}^{-1}(r, 2N)B_{2N}J_{12}\phi_{1\sim 2}^{-1}(r, 2N)^T \right],$$

$$r \leq 2N,$$

can be separated for $r \leq N$ as:

$$\begin{aligned}\phi_{1\sim 4}^{-1}(r, r) &= \phi_1^{-1}(r, r) - \left\{ \left[\phi_2^{-1}(N, 1)B_N^TJ_{22}\phi_1^{-1}(r, N)^T \right]^T \right\} \\ &\quad \cdot B_{2N}J_{12} \cdot \left\{ \left[\phi_2^{-1}(N, 1) + B_N^TJ_{22}\phi_1^{-1}(r, N)^T \right] \right\},\end{aligned}$$

$$\begin{aligned}\phi_{1\sim 4}^{-1}(r, N, r + N) &= \phi_2^{-1}(r, r) \\ &\quad - \left\{ \left[\phi_2^{-1}(N, r) + \phi_2^{-1}(N, 1)B_N^TJ_{21}\phi_2^{-1}(1, r) \right]^T \right\} \cdot B_{2N}J_{12} \\ &\quad \cdot \left\{ \left[\phi_2^{-1}(N, r) + \phi_2^{-1}(N, 1)B_N^TJ_{21}\phi_2^{-1}(1, r) \right] \right\}.\end{aligned}$$

Thus, the modifications to the diagonal entries can be written as just a function of the first block row and last block column from the individual sub-matrices, using the matrix map framework introduced in (A4) for $r \leq N$,

$$\begin{aligned}\phi_{1\sim 4}^{-1}(r, r) &= \phi_2^{-1}(r, r) \\ &\quad - \left\{ \left[M_{1;1}\phi_1^{-1}(1, r) + M_{2;1}\phi_1^{-1}(r, N)^T \right]^T \right\} \cdot B_{2N}J_{12} \\ &\quad \cdot \left\{ \left[M_{1;1}\phi_1^{-1}(1, r) + M_{2;1}\phi_1^{-1}(r, N)^T \right] \right\}, \\ \phi_{1\sim 4}^{-1}(r + N, r + N) &= \phi_2^{-1}(r, r) \\ &\quad - \left\{ \left[M_{1;2}\phi_2^{-1}(1, r) + M_{2;2}\phi_2^{-1}(r, N)^T \right]^T \right\} \cdot B_{2N}J_{12} \\ &\quad \cdot \left\{ \left[M_{1;2}\phi_2^{-1}(1, r) + M_{2;2}\phi_2^{-1}(r, N)^T \right] \right\}.\end{aligned}$$

Here, the matrix maps are assumed to have been updated based upon the formation of the combined inverses $\phi_{1\sim 2}^{-1}$ and $\phi_{3\sim 4}^{-1}$. Therefore, we can begin to formulate the recursive framework for updating the matrix maps to represent the effect of each combining step.

3. Update scheme for parallel inversion and the distributed compact representation

The procedure begins with each division of the problem being assigned to one of p available CPUs. In addition, all of the $p-1$ bridge matrices are made available to each of the CPUs. After the compact representation for each inverse has been found independently, the combining process begins. Three reference positions are defined for the formation of a combined inverse $\phi_{i\sim j}^{-1}$: the “start” position $[\text{st}] = i$, the “stop” position $[\text{sp}] = j$, and the bridge position $[\text{bp}] = \lceil (j-i)/2 \rceil$. Due to the fact that a CPU t will only be involved in the formulation of a combined inverse when $[\text{st}] \leq t \leq [\text{sp}]$ all combining stages on the same level (see Fig. 1) can be performed concurrently. When forming a combined inverse $\phi_{i\sim j}^{-1}$, each CPU $[\text{st}] \leq t \leq [\text{sp}]$ will first need to form the adjustment matrix for the combining step. Assuming a bridge matrix B_k , we begin by constructing four “corner blocks.” If the upper combined inverse is assumed to

have N_u blocks and the lower to have N_l , the two matrices need from the upper combined inverse are $[\text{UR}] = \phi_{[\text{st}] \sim [\text{bp}]}^{-1}(1, N_u)$ and $[\text{LR}] = \phi_{[\text{st}] \sim [\text{bp}]}^{-1}(N_u, N_u)$, with the two matrices from the lower being: $[\text{UL}] = \phi_{[\text{bp}+1] \sim [\text{sp}]}^{-1}(1, 1)$ and $[\text{LL}] = \phi_{[\text{bp}+1] \sim [\text{sp}]}^{-1}(N_l, 1)$. These matrices can be generated by the appropriate CPU through their respective matrix maps (recall the example shown in Fig. 6). Specifically, the CPUs corresponding to the $[\text{st}]$, $[\text{bp}]$, $[\text{bp} + 1]$, and $[\text{sp}]$ sub-matrices govern the required information. The adjustment matrix for the combining stage can then be formed

$$J = \begin{pmatrix} I & -[\text{UL}]B_k^T \\ -[\text{LR}]B_k & I \end{pmatrix}^{-1}.$$

After the adjustment matrix has been calculated the process of updating the matrix maps can begin. For any combining stage, the cross maps each CPU t must be updated first,

if $(t < [\text{bp}])$ then

$$\begin{aligned} C_1 &\leftarrow C_1 - M_{3;t}^T(B_k J_{12})M_{3;t+1}; \\ C_2 &\leftarrow C_2 - M_{3;t}^T(B_k J_{12})M_{4;t+1}; \\ C_3 &\leftarrow C_3 - M_{4;t}^T(B_k J_{12})M_{3;t+1}; \\ C_4 &\leftarrow C_4 - M_{4;t}^T(B_k J_{12})M_{4;t+1}; \\ \text{elseif } (t == [\text{bp}]) &\text{ then} \\ C_1 &\leftarrow C_1 - M_{3;t}^T(B_k J_{11})M_{1;t+1}; \\ C_2 &\leftarrow C_2 - M_{3;t}^T(B_k J_{11})M_{2;t+1}; \\ C_3 &\leftarrow C_3 - M_{4;t}^T(B_k J_{11})M_{1;t+1}; \\ C_4 &\leftarrow C_4 - M_{4;t}^T(B_k J_{11})M_{2;t+1}; \\ \text{elseif } (t < [\text{sp}]) &\text{ then} \\ C_1 &\leftarrow C_1 - M_{1;t}^T(B_k^T J_{21})M_{1;t+1}; \\ C_2 &\leftarrow C_2 - M_{1;t}^T(B_k^T J_{21})M_{2;t+1}; \\ C_3 &\leftarrow C_3 - M_{2;t}^T(B_k^T J_{21})M_{1;t+1}; \\ C_4 &\leftarrow C_4 - M_{2;t}^T(B_k^T J_{21})M_{2;t+1}. \end{aligned} \quad (\text{A8})$$

Notice that the cross maps for CPU t are dependent on information from its neighboring CPU $t + 1$. This information must be transmitted and made available before the cross updates can be performed. Next, updates to the remaining eight matrix maps can be separated into two categories. The updates to the matrix maps for the upper sub-matrices ($t \leq [\text{bp}]$), are summarized below.

$$\begin{aligned} M_{5;t} &\leftarrow M_{5;t} - M_{3;t}^T(B_k J_{12})M_{3;t}; \\ M_{6;t} &\leftarrow M_{6;t} - M_{3;t}^T(B_k J_{12})M_{4;t}; \\ M_{7;t} &\leftarrow M_{7;t} - M_{4;t}^T(B_k J_{12})M_{3;t}; \\ M_{8;t} &\leftarrow M_{8;t} - M_{4;t}^T(B_k J_{12})M_{4;t}; \\ M_{1;t} &\leftarrow M_{1;t} + ([\text{UR}]B_k J_{12})M_{3;t}; \\ M_{2;t} &\leftarrow M_{2;t} + ([\text{UR}]B_k J_{12})M_{4;t}; \\ M_{3;t} &\leftarrow ([\text{LL}]^T B_k^T J_{22})M_{4;t}; \\ M_{4;t} &\leftarrow ([\text{LL}]^T B_k^T J_{22})M_{4;t}. \end{aligned} \quad (\text{A9})$$

The updates to the matrix maps for the lower sub-matrices ($t > [\text{bp}]$), will be

$$\begin{aligned} M_{5;t} &\leftarrow M_{5;t} - M_{1;t}^T(B_k^T J_{21})M_{1;t}; \\ M_{6;t} &\leftarrow M_{6;t} - M_{1;t}^T(B_k^T J_{21})M_{2;t}; \\ M_{7;t} &\leftarrow M_{7;t} - M_{2;t}^T(B_k^T J_{21})M_{1;t}; \\ M_{8;t} &\leftarrow M_{8;t} - M_{2;t}^T(B_k^T J_{21})M_{2;t}; \\ M_{3;t} &\leftarrow M_{3;t} + ([\text{LL}]^T B_k^T J_{21})M_{1;t}; \\ M_{4;t} &\leftarrow M_{4;t} + ([\text{LL}]^T B_k^T J_{21})M_{2;t}; \\ M_{1;t} &\leftarrow ([\text{UR}]B_k J_{11})M_{1;t}; \\ M_{2;t} &\leftarrow ([\text{UR}]B_k J_{11})M_{2;t}. \end{aligned} \quad (\text{A10})$$

The above procedure, shown in (A8)–(A10), for modifying the matrix maps can be recursively repeated for each of the combining stages beginning with the lowest level of combining the individual sub-matrix inverses.

On completion the maps can then be used to generate the block tridiagonal entries of G^R . This subsequently allows for the computation of the generator sequences for G^R , via the relationships shown in (4), in a purely distributed fashion. The time complexity of the algorithm presented is $O(N_x^3 N_y / p + N_x^3 \log p)$, with memory consumption $O(N_x^2 N_y / p + N_x^2)$. The first term ($N_x^3 N_y / p$) in the computational complexity arises from the embarrassingly parallel nature of both determining the generator sequences and applying the matrix maps to update the block tridiagonal portion of the inverse. The second term ($N_x^3 \log p$) is dependent on the number of levels needed to gather combining information for p sub-matrix inverses. Similarly, the first term in the memory complexity is due to the generator sequences and diagonal blocks, and the second represents the memory required for the matrix maps of each sub-problem governed.

ACKNOWLEDGMENTS

This work was partially supported by NSF Grant No. EEC-0228390, by NSF PetaApps Grant No. 0749140, by NSF through TeraGrid resources provided by the National Institute for Computational Sciences, and by DOE resources provided by the National Center for Computational Sciences under Contract No. DE-AC05-00OR22725.

¹A. Svizhenko, M. P. Anantram, T. R. Govindan, B. Biegel, and R. Venugopal, *J. Appl. Phys.* **91**(4), 2343 (2002).

²S. Cauley, J. Jain, C.-K. Koh, and V. Balakrishnan, *J. Appl. Phys.* **101**, 123715 (2007).

³S. Li, S. Ahmed, G. Klimeck, and E. Darve, *J. Comput. Phys.* **227**(22), 9408 (2008).

⁴M. Stadele, R. Tuttle, and K. Hess, *J. Appl. Phys.* **89**(1), 348 (2001).

⁵T. A. Davis, *ACM Trans. Math. Software* **30**(2), 196 (2004).

⁶J. W. Demmel, S. C. Eisenstat, J. R. Gilbert, X. S. Li, and J. W. H. Liu, *SIAM J. Matrix Analysis and Applications* **20**(3), 720 (1999).

⁷T. B. Boykin, M. Luisier, and G. Klimeck, *Phys. Rev. B* **77**, 165318 (2008).

⁸K. Bowden, *Int. J. Gen. Syst.* **15**, 185 (1989).

⁹B. Bukhberger and G. A. Emel'yanenko, *Comp. Math. Math. Phys.* **13**, 10 (1973).

¹⁰E. M. Godfrin, *J. Phys.: Condens. Matter* **3**, 7843 (1991).

¹¹T. Oohashi, *TRU Mathematics* **14**, 39 (1978).

¹²T. Torii, *Technology Reports of the Osaka University* **16**, 403 (1966).

- ¹³G. Meurant, *SIAM Journal on Matrix Analysis and Applications* **13**(3), 707 (1992).
- ¹⁴E. Asplund, *Mathematica Scandinavica* **7**, 57 (1959).
- ¹⁵S. O. Asplund, *Mathematica Scandinavica* **7**, 49 (1959).
- ¹⁶R. Bevilacqua, B. Codenotti, and F. Romani, *Linear Algebra Appl.* **104**, 39 (1988).
- ¹⁷R. Nabben, *SIAM J. Matrix Analysis and Applications* **20**(3), 820 (1999).
- ¹⁸F. Romani, *Linear Algebra Appl.* **80**, 131 (1986).
- ¹⁹P. Rozsa, *Integral Equations and Operator Theory* **50**, 82 (1987).
- ²⁰P. Rozsa, R. Bevilacqua, P. Favati, and F. Romani, *Operator Theory: Advances and Applications* **40**, 447 (1989).
- ²¹P. Rozsa, *Colloquia Mathematica Societatis Janos Bolyai* **50**, 22 (1986).
- ²²P. Concus and G. Meurant, *BIT* **26**, 493 (1986).
- ²³S. Demko, W. F. Moss, and P. W. Smith, *Math. Comput.* **43**, 491 (1984).
- ²⁴P. Concus, G. H. Golub, and G. Meurant, *SIAM J. Sci. Comput.* **6**, 220 (1985).
- ²⁵M. Luisier and G. Klimeck, *Phys. Rev. B* **80**, 155430 (2009).
- ²⁶J. Jain, S. Cauley, H. Li, C.-K. Koh, and V. Balakrishnan, Purdue ECE Technical Report 1358 (2007).
- ²⁷K. Rim, S. Narasimha, M. Longstreet, A. Mocuta, and J. Cai, *IEDM Tech. Dig.* 43 (2002).
- ²⁸M. Luisier, *Appl. Phys. Lett.* **98**, 032111 (2011).