SPARSE MATRIX-MATRIX PRODUCTS EXECUTED THROUGH COLORING

MICHAEL MCCOURT *, BARRY SMITH *, AND HONG ZHANG *

Abstract. Sparse matrix-matrix products appear in multigrid solvers among other applications. Some implementations of these products require the inner product of two sparse vectors. In this paper, we propose a new algorithm for computing sparse matrix-matrix products by exploiting their nonzero structure through the process of graph coloring. We prove the validity of this technique in general and demonstrate its viability for examples including multigrid methods used to solve boundary value problems as well as matrix products appearing in unstructured applications.

Key words. sparse matrix product, coloring

AMS subject classifications. 65F50, 65F30

1. Introduction. Matrix-matrix multiplication is a fundamental linear algebra operation [10]. The operation of concern in this work is

$$\mathsf{C} = \mathsf{A}\mathsf{B}^T,$$

which arises in algebraic multigrid. This operation is well defined when $\mathsf{A} \in \mathbb{R}^{m\times r}$ and $\mathsf{B} \in \mathbb{R}^{n\times r}$ and produces $\mathsf{C} \in \mathbb{R}^{m\times n}$. Our results are equally applicable to complex-valued matrices, but we omit that discussion for simplicity.

The multiplication operation is defined such that the jth column of the ith row of C (denoted as $\mathsf{C}(i,j)$) is calculated as

$$\mathsf{C}(i,j) = \sum_{k=1}^{r} \mathsf{A}(i,k)\mathsf{B}^T(k,j).$$

This can also be written as the *inner product* of the ith row of A, denoted as $\mathsf{A}(i,:)$, and the jth column of B^T, denoted as $\mathsf{B}^T(:,j)$:

$$\mathsf{C}(i,j) = \mathsf{A}(i,:)\mathsf{B}^T(:,j). \tag{1.1}$$

Several mechanisms exist for computing the matrix-matrix product, each of which is preferable in certain settings. For instance, C can be constructed all at once with a sum of r outer products (rank-1 matrices),

$$\mathsf{C} = \sum_{k=1}^{r} \mathsf{A}(:,k)\mathsf{B}^T(k,:). \tag{1.2}$$

For an approach between using (1.1) to compute one value at a time and using (1.2) to compute C all at once, we can compute one row or column at a time or compute blocks of C using blocks of A and B^T.

While these techniques all yield the same result, they may not be equally preferable for actual computation because of the different forms in which a matrix may be stored. Many sparse matrix algorithms have memory access outpace floating-point operations as the computational bottleneck; this has arisen in uniprocessor settings

*Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439. Emails: `(mccomic, bsmith, hzhang)@mcs.anl.gov`

[20], parallel algorithms [4] and in developing tools for GPU computation [6]. In this paper, we begin with conventional matrix storage: compressed sparse row format (CSR) for sparse matrices and Fortran-style column-major array storage for dense matrices [3] and then transform the storage format to improve efficiency. A matrix $\mathsf{A} \in \mathbb{R}^{m\times n}$ stored in CSR format requires only $2nnz + m$ storage, where nnz is the number of nonzeros in A; accessing consecutive values in a row of a CSR matrix is very efficient, whereas accessing consecutive values in a column is not. As a result, algorithms which most efficiently use CSR matrices involve entire rows of the matrix to maximize the amount of computation before another call to memory must occur.

One approach to computing AB^T would be to compute B^T from B and store it as a CSR matrix, but this reorganization requires a significant amount of memory traffic. The availability of the columns of B^T would seem to indicate that the inner product computation (1.1) is the preferred method; however, computing the inner product of sparse vectors can be inefficient operation (see Section 1.1 or [11]). Our goal is to produce a data structure and algorithm which efficiently computes the sparse matrix-matrix product $\mathsf{C} = \mathsf{AB}^T$ using inner products.

In the rest of this section we introduce sparse inner products and matrix coloring. In Section 2 we analyze matrix coloring applied to the sparse matrix product $\mathsf{C} = \mathsf{AB}^T$, which allows us to instead compute C by evaluating the inner product of sparse and dense vectors. In Section 3 we propose algorithms for computing matrix products with matrix colorings and consider some practical implementation issues. Numerical results are presented in Section 4, where some favorable results in comparison to other matrix multiplication strategies. These experiments are conducted on the matrix products appearing in multigrid solvers for three dimensional PDEs from PFLOTRAN [15], from the PETSc library distribution [3], RBF network [18] and the University Florida the University of Florida Sparse Matrix Collection [7]. We conclude the paper in Section 5 with a brief discussion of future work.

1.1. Sparse inner products. When matrices A and B are stored in compressed sparse row format, computing $\mathsf{C} = \mathsf{AB}^T$ requires the inner product between sparse vectors; this section discusses the potential inefficiency of sparse inner products. Algorithm 1 shows how inner products between sparse vectors in $\mathbb{R}^n$ may be computed.

Algorithm 1 Sparse-Sparse Inner Product of $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$; $\boldsymbol{x}, \boldsymbol{y}$ have n_x, n_y nonzeros, respectively.

1: $i = 1$; $j = 1$; $\boldsymbol{x}^T\boldsymbol{y} = 0.0$
2: **while** $(i \le n_x$ and $j \le n_y)$ **do**
3: if ($\text{index}(\boldsymbol{x}, i) < \text{index}(\boldsymbol{y}, j)$) { $i = i + 1$; }
4: else if ($\text{index}(\boldsymbol{x}, i) > \text{index}(\boldsymbol{y}, j)$) { $j = j + 1$; }
5: else { $\boldsymbol{x}^T\boldsymbol{y} = \boldsymbol{x}^T\boldsymbol{y} + \boldsymbol{x}(\text{index}(\boldsymbol{x}, i))\boldsymbol{y}(\text{index}(\boldsymbol{y}, j))$; $i = i + 1$; $j = j + 1$;}
6: **end while**

Because $\boldsymbol{x}$ and $\boldsymbol{y}$ are stored in compressed form, only the n_x and n_y nonzero values in each vector respectively are stored, along with the rows to which they belong. The function "index" accepts a vector $\boldsymbol{x}$ and an integer $1 \le i \le n_x$ and returns the index of the ith nonzero; the ith nonzero is then accessed as $\boldsymbol{x}(\text{index}(\boldsymbol{x}, i))$.

Algorithm 1 iterates through the nonzero entries in each vector until finding the collisions between their indices (indices where $\boldsymbol{x}(\text{index}(\boldsymbol{x}, i))\boldsymbol{y}(\text{index}(\boldsymbol{y}, j)) \neq 0$); each collision is then accumulated to evaluate the inner product. While the number of floating-point operations (flops) for this algorithm is only proportional to the number

of collisions between nonzeros, the required memory access is inefficient. The indices of both vectors must be traversed completely, requiring $O(n_x + n_y)$ memory accesses in a while-loop.

This poor ratio of computation to memory access is one motivation behind our work. Contrast Algorithm 1 with the inner product of a sparse and dense vector in Algorithm 2.

Algorithm 2 Sparse-Dense Inner Product of $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$; $\boldsymbol{x}$ has n_x nonzeros and $\boldsymbol{y}$ is dense.

1: $\boldsymbol{x}^T\boldsymbol{y} = 0.0$
2: **for** $(i = 1, \dots, n_x)$ **do**
3: $\quad \boldsymbol{x}^T\boldsymbol{y} = \boldsymbol{x}^T\boldsymbol{y} + \boldsymbol{x}(\text{index}(\boldsymbol{x}, i))\boldsymbol{y}(\text{index}(\boldsymbol{x}, i))$
4: **end for**

Algorithm 2 performs $O(n_x)$ flops, because even if some of the values in $\boldsymbol{y}$ are zero, they are all treated as nonzeros. This is greater than the sparse-sparse inner product but with roughly the same level of memory accesses: $O(2n_x)$ in a for-loop. Essentially, a sparse-dense inner product would do more flops per memory access. Our interest is not in this small trade-off but in compressing multiple sparse vectors into a single, dense vector; this significantly increases the ratio of flops to memory accesses, without (we hope) introducing a harrowing number of zero-valued nonzeros into the computation. If this can be done effectively, then the inefficient sparse-sparse inner products used to compute sparse matrix products can be replaced with a reduced number of more efficient sparse-dense inner products. Our mechanism for doing this is described in Section 1.2.

1.2. Matrix coloring. Matrix coloring is related to graph coloring, an important topic in graph theory [14]. The original use of graph coloring in matrices was in approximating the Jacobian of a nonlinear system using fewer function evaluations by exploiting sparsity [5]; a more comprehensive paper explaining the role of graph coloring for Jacobians is [9]. We are not interested in a graph-theoretic understanding, only in the specific use of graph theory to color matrices, so we introduce only terms specifically relevant to our usage.

DEFINITION 1.1. *Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^n$. The vectors $\boldsymbol{u}$ and $\boldsymbol{v}$ are* structurally orthogonal *if $|\boldsymbol{u}|^T|\boldsymbol{v}| = 0$, where $|\boldsymbol{u}|$ is the vector of the absolute value of all the elements of $\boldsymbol{u}$. A set of vectors $\{\boldsymbol{u}_1, ..., \boldsymbol{u}_n\}$ is called* structurally orthogonal *if $\boldsymbol{u}_i$ and $\boldsymbol{u}_j$ are structurally orthogonal for $1 \leq i, j \leq n$.*

Under this definition, not all orthogonal vectors are structurally orthogonal; for example,

$$\boldsymbol{u} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \qquad \boldsymbol{v} = \begin{pmatrix} 1 \\ -1 \end{pmatrix},$$

so $\boldsymbol{u}^T\boldsymbol{v} = 0$ but $|\boldsymbol{u}|^T|\boldsymbol{v}| = 2$. Also, all vectors are structurally orthogonal to a vector of all zeros.

LEMMA 1.2. *Let $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{R}^n$ be structurally orthogonal, and denote $\boldsymbol{u}(k)$ as the kth element in $\boldsymbol{u}$. Then for $1 \leq k \leq n$, either $\boldsymbol{u}(k) = 0$ or $\boldsymbol{v}(k) = 0$, or both.*

Proof. The proof follows trivially from the definition. ☐

DEFINITION 1.3. *Let $\mathsf{C} \in \mathbb{R}^{m \times n}$. An* orthogonal set *is a set of q indices $\boldsymbol{\ell} = \{\ell^1, \ell^2, \dots, \ell^q\}$ for which $\mathsf{C}(:, \ell^i)$ and $\mathsf{C}(:, \ell^j)$ are structurally orthogonal when $1 \leq$*

$i, j \leq q$ and $i \neq j$. We also define a set containing only one index $\boldsymbol{\ell} = \{\ell^1\}$ to be an orthogonal set.

Definition 1.4. *Let $\mathsf{C} \in \mathbb{R}^{m \times n}$. A* matrix coloring *$c = \{\boldsymbol{\ell}_1, \boldsymbol{\ell}_2, \ldots, \boldsymbol{\ell}_p\}$ is a collection of index sets such that, for $1 \leq k \leq n$, the index k appears in exactly one index set. We say that c is a* valid matrix coloring *of C if each index set in c is an orthogonal set of C. We refer to an orthogonal set that is part of a coloring as a* color.

Because we require $1 \leq k \leq n$ to appear in exactly one orthogonal set, every column of C appears in exactly one color. Each orthogonal set in the coloring contains a set of indices corresponding to columns of C which are structurally orthogonal. The term coloring mimics the graph coloring concept of grouping nonadjacent vertices using the same color; here we are grouping structurally orthogonal columns of a matrix using the same orthogonal set.

Recall that we allow for the possibility of there being only one column in a color; for instance, a column with no zero values in it must exist in its own color. This in turn guarantees that every matrix has a coloring, since every matrix $\mathsf{C} \in \mathbb{R}^{m \times n}$ must have at least the trivial coloring $\{\{1\}, \{2\}, \ldots, \{n\}\}$. If the matrix C were totally dense, with no structural zeros at all, this would be the only coloring. Our focus, however, is on the coloring of very sparse matrices.

A matrix may have more than one valid coloring; the 2×2 identity

$$\mathsf{I}_2 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$$

has two colorings: $c_1 = \{\{1, 2\}\}$ and $c_2 = \{\{1\}, \{2\}\}$. We are not concerned with how the coloring of a matrix is determined but rather how we can use it to efficiently compute matrix-matrix products. Therefore, we refer readers to [9] for an introduction to how colorings are found.

Definition 1.5. *Let $\mathsf{C} \in \mathbb{R}^{m \times n}$ and $c = \{\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_p\}$ be a valid matrix coloring for C.* Applying the coloring *c to C produces a matrix C_{Dense} that is at least as dense as the matrix C. If c has p colors in it, the matrix $\mathsf{C}_{Dense} \in \mathbb{R}^{m \times p}$. The kth column of the matrix C_{Dense} is created by combining the columns of C from the orthogonal set $\boldsymbol{\ell}_k = \{\ell_k^1, \ldots, \ell_k^{q_k}\}$. Specifically,*

$$\mathsf{C}_{Dense}(j, k) = \begin{cases} \mathsf{C}(j, \ell_k^r), & \text{if} \quad \mathsf{C}(j, \ell_k^r) \neq 0, \\ 0, & \text{if} \quad \sum_{i=1}^{q_k} |\mathsf{C}(j, \ell_k^i)| = 0, \end{cases}$$

for $1 \leq j \leq m$. The matrix resulting from applying the coloring is referred to as a compressed matrix.

We note that it is also possible to define $\mathsf{C}_{Dense} = \mathsf{C}\mathsf{D}$ for a sparse matrix $\mathsf{D} \in \mathbb{R}^{m \times p}$ such that all nonzeros satisfy $\mathsf{D}(j, k) = 1$ if $j \in \boldsymbol{\ell}_k$. Such a definition is more compact, but the definition above is useful for the proofs in this article.

Theorem 1.6. *Let $\mathsf{C} \in \mathbb{R}^{m \times n}$ and $c = \{\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_p\}$ be a valid matrix coloring for C. The matrix C_{Dense} created by applying c to C is unique.*

Proof. See appendix.

Traditionally, matrix colorings have been used for the accelerated computation of finite-difference Jacobians for the purpose of preconditioning nonlinear solvers [5, 9]. In that setting, structurally orthogonal columns were grouped together to allow for simultaneous function evaluation. This minimizes the number of function calls required to approximate the Jacobian without ignoring any nonzero values. We use this same concept in Section 2 to accelerate the operation $\mathsf{C} = \mathsf{A}\mathsf{B}^T$.

2. Analysis of Coloring for Sparse Matrix-Matrix Products. Our goal in this section is to exploit the graph coloring concept, described in Section 1.2, to more efficiently compute inner products in a sparse-sparse matrix product. We are interested in the two expressions

$$\mathsf{C} = \mathsf{AB}^T \text{ or} \tag{2.1a}$$

$$\mathsf{C} = \mathsf{RAR}^T, \tag{2.1b}$$

where matrices $\mathsf{A}, \mathsf{B}, \mathsf{R}$ are stored in CSR format and are not necessarily square. We analyze only (2.1a) in this section because (2.1b) can be described by using this product.

The sparse-sparse matrix product (2.1a) can be computed with an analagous sparse-dense matrix product; doing so reduces memory access overhead and improves computational efficiency. To achieve this improved performance, we restructure the sparse matrix B^T into a dense matrix B^T_{Dense} using the matrix coloring ideas introduced in Section 1.2. This allows us to compute

$$\mathsf{C}_{Dense} = \mathsf{AB}^T_{Dense}, \tag{2.2}$$

such that the nonzero entries in C_{Dense} are also the nonzero entries needed to form C. After computing (2.2), C_{Dense} must be reorganized into C, which is done by using a process similar to the compression of B^T into B^T_{Dense}. Since this is not fundamental to the analysis, this process is discussed in Section 3.1.

We must ask the question: How can we form a dense B^T_{Dense} such that C_{Dense} and C have the same nonzero values? Our approach uses a coloring of the matrix C to determine which columns of C are structurally orthogonal and can be computed simultaneously. This allows for the compression of those multiple columns associated with each color of B^T into a single column of B^T_{Dense}. After C_{Dense} has been computed, this single dense column is decompressed to fill all the columns of C corresponding to the columns of B^T, which were earlier compressed to created B^T_{Dense}.

Although the dense matrix C_{Dense} contains the same nonzero values as C, the dense matrix, B^T_{Dense} need not contain all the nonzero values of B^T. This special case may occur when A has columns with only zero values, but we prove in Corollary 2.7 that whenever A has no zero columns, any coloring of C *is a valid coloring* for B^T.

2.1. Motivating Example. We consider a small example to demonstrate the coloring concept before we study its validity in general sparse-sparse matrix products. Consider the product of two "sparse" matrices,

$$\underbrace{\begin{pmatrix} 1 & 6 & 3 & & & 12 \\ & 4 & & & & 12 \\ 1 & 2 & 3 & & & 36 \\ & & 12 & 28 & 15 & 24 \\ & & & 20 & & 66 \\ & & & & & 36 \end{pmatrix}}_{\mathsf{C}} = \underbrace{\begin{pmatrix} 1 & 2 & & & & \\ & 2 & & & & \\ 1 & & & & & 6 \\ & & 3 & 4 & & \\ & & & & 5 & 6 \\ & & & & & 6 \end{pmatrix}}_{\mathsf{A}} \underbrace{\begin{pmatrix} 1 & 2 & 3 & & & \\ & 2 & & & & 6 \\ & & & 4 & 5 & \\ & & 3 & 4 & & 6 \\ & & & 4 & & 6 \\ & & & & & 6 \end{pmatrix}}_{\mathsf{B}^T}.$$

The structure of C admits the coloring

$$c = \{\ell_1, \ell_2, \ell_3, \ell_4\} = \{\{1,4\}, \{2,5\}, \{3\}, \{6\}\}$$

because columns 1 and 4 are structurally orthogonal, as are columns 2 and 5. Notice that despite the structural orthogonality of the third and fifth columns of B^T, those columns cannot be combined in the same color in C. Using the coloring c, we can compress B^T into B^T_{Dense} and compute C_{Dense},

$$\underbrace{\begin{pmatrix} 1 & 6 & 3 & 12 \\ & 4 & & 12 \\ 1 & 2 & 3 & 36 \\ 28 & 15 & 12 & 24 \\ 20 & & & 66 \\ & & & 36 \end{pmatrix}}_{\mathsf{C}_{Dense}} = \underbrace{\begin{pmatrix} 1 & 2 & & & & \\ & 2 & & & & \\ 1 & & & & & 6 \\ & & 3 & 4 & & \\ & & & & 5 & 6 \\ & & & & & 6 \end{pmatrix}}_{\mathsf{A}} \underbrace{\begin{pmatrix} 1 & 2 & 3 & \\ & 2 & & 6 \\ 4 & 5 & & \\ 4 & & 3 & 6 \\ 4 & & & 6 \\ & & & 6 \end{pmatrix}}_{\mathsf{B}^T_{Dense}}.$$

To decompress C_{Dense} into C, we must store both the columns that formed the coloring and the rows of C associated with each of those columns. In the example above, the first column of C_{Dense} is composed of columns 1 and 4 from C, and the fourth and fifth rows belong to column 4 of C. This process is discussed in Section 3.1.

2.2. Proof of the Validity of Matrix Coloring. The example in the previous subsection suggests that a valid coloring of C can be used to compress B^T; indeed, whenever a coloring for C is also valid for B^T, Theorem 1.6 says that B^T_{Dense} is unique and thus $\mathsf{C}_{Dense} = \mathsf{A}\mathsf{B}^T_{Dense}$ is the matrix of interest. In this section, we prove that even when c is not a valid coloring for B^T, if it is a valid coloring for C we can still use it for our matrix product. Before we can complete such a proof, we need to create a new device: an auxiliary matrix $\hat{\mathsf{B}}^T$ for which c is a valid coloring.

DEFINITION 2.1. *Let $\mathsf{B}^T \in \mathbb{R}^{r\times n}$ and c be a valid coloring for some matrix with n columns; c need not be a valid coloring for B^T. A* vacated matrix *$\hat{\mathsf{B}}^T \in \mathbb{R}^{r\times n}$ is a matrix whose nonzero values all coincide with nonzeros from B^T but for which c is a valid coloring.*

Under this definition, no nonzeros can be introduced in vacating B^T to $\hat{\mathsf{B}}^T$; values can only be zeroed out to produce a matrix with the appropriate structure to apply the coloring. If c is a valid coloring of B^T, then one trivial vacation of B^T would be to remove no nonzeros. Another trivial vacation of B^T that would validate any coloring would be to zero out every value in the matrix.

DEFINITION 2.2. *Let $\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_q \in \mathbb{R}^r$. We define the* conflicted index set *of $\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_q\}$ as the set of indices*

$$\Gamma(\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_q\}) = \{\, \gamma \in \{1, 2, \ldots, r\} \mid \boldsymbol{u}_i(\gamma)\boldsymbol{u}_j(\gamma) \neq 0, \textit{ for some } 1 \leq i, j \leq q,\ i \neq j\}.$$

We define the conflicted index set of a set of one vector as empty: $\Gamma(\{\boldsymbol{u}\}) = \emptyset$.

The term *conflicted index set* refers to the fact that these rows are preventing the set of vectors $\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_q\}$ from being structurally orthogonal. Were they structurally orthogonal, then the conflicted index set would be empty.

LEMMA 2.3. *Let $\boldsymbol{u}_1, \boldsymbol{u}_2, \ldots, \boldsymbol{u}_q \in \mathbb{R}^r$, and $\Gamma(\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_q\})$ be the associated conflicted index set. The set of vectors $\{\hat{\boldsymbol{u}}_1, \hat{\boldsymbol{u}}_2, \ldots, \hat{\boldsymbol{u}}_q\}$ defined as*

$$\hat{\boldsymbol{u}}_i(\gamma) = \begin{cases} 0 & \gamma \in \Gamma(\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_q\}) \\ \boldsymbol{u}_i(\gamma) & \textit{else} \end{cases}$$

is structurally orthogonal.

Proof. See appendix.

This lemma allows us to take any set of vectors and replace certain nonzero values with zeros to make them structurally orthogonal. We apply this concept to the columns of B^T involved in each color of c in order to create a matrix $\hat{\mathsf{B}}^T$ for which c is a valid coloring.

DEFINITION 2.4. *Let* $\mathsf{B}^T \in \mathbb{R}^{r\times n}$, $c = \{\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_p\}$ *be a coloring for some matrix with* n *columns, and let* $\mathsf{B}^T(:, \boldsymbol{\ell}_k)$ *denote the set of columns associated with the color* $\boldsymbol{\ell}_k = \{\ell_k^1, \ldots, \ell_k^{q_k}\}$. *The* minimally vacated matrix $\hat{\mathsf{B}}_c^T \in \mathbb{R}^{r\times n}$ *is the vacated matrix generated from* B^T *in the following way:*

$$\hat{\mathsf{B}}_c^T(\gamma, \ell_k^j) = \begin{cases} 0, & \gamma \in \Gamma(\mathsf{B}^T(:, \boldsymbol{\ell}_k)) \\ \mathsf{B}^T(\gamma, \ell_k^j), & else \end{cases}, \qquad 1 \le j \le q_k, \quad 1 \le k \le p.$$

LEMMA 2.5. *Let* $\mathsf{B}^T \in \mathbb{R}^{r\times n}$, $c = \{\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_p\}$ *be a valid coloring for some matrix with* n *columns. The coloring* c *is valid for the minimally vacated matrix* $\hat{\mathsf{B}}_c^T$.

Proof. Using Lemma 2.3, we know that the columns in $\hat{\mathsf{B}}_c^T(:, \boldsymbol{\ell}_k)$ are structurally orthogonal for $1 \le k \le p$, and therefore c is a valid coloring of $\hat{\mathsf{B}}_c^T$. □

A minimally vacated matrix is a vacated matrix where nonzeros have been removed specifically to validate the given coloring; that is, only indices appearing in the conflicted index set for each color are zeroed-out. Recall that for the product $\mathsf{C} = \mathsf{A}\mathsf{B}^T$, our goal is to produce a dense matrix B_{Dense}^T from the matrix B^T by applying the coloring c to B^T. As stated earlier, c may not be a valid coloring for B^T, in which case we substitute the minimally vacated matrix $\hat{\mathsf{B}}_c^T$ in place of B^T. Lemma 2.5 proves that c is an acceptable coloring for this minimally vacated matrix.

The term "minimally" vacated is used to suggest that no values needed for the computation $\mathsf{A}\mathsf{B}^T$ are lost, in contrast to another vacation, such as removal of all the nonzeros of the matrix, which would validate the coloring but probably lose needed information. More important than the validity of the coloring is the validity of the equation $\mathsf{C} = \mathsf{A}\hat{\mathsf{B}}_c^T$. If this equality does not hold, then some values in C would be altered by substituting the minimally vacated $\hat{\mathsf{B}}_c^T$ for B^T, which is unacceptable. Theorem 2.6 addresses this concern.

THEOREM 2.6. *Let* $\mathsf{A} \in \mathbb{R}^{m\times r}$, $\mathsf{B} \in \mathbb{R}^{n\times r}$ *and* $c = \{\boldsymbol{\ell}_1, \ldots, \boldsymbol{\ell}_p\}$ *be a coloring of* $\mathsf{C} = \mathsf{A}\mathsf{B}^T$. *If* $\hat{\mathsf{B}}_c^T$ *is the minimally vacated matrix generated by applying* c *to the matrix* B^T, *then the equality*

$$\mathsf{A}\mathsf{B}^T = \mathsf{A}\hat{\mathsf{B}}_c^T \tag{2.3}$$

holds.

Proof. We denote $\hat{\mathsf{C}} = \mathsf{A}\hat{\mathsf{B}}_c^T$ and prove that $\mathsf{C} = \hat{\mathsf{C}}$ by proving that all the nonzeros of $\hat{\mathsf{C}}$ match the nonzeros of C. By definition, each column of C belongs to exactly one color, so let us denote as $\boldsymbol{\ell}_{p_j} = \{\ell_{p_j}^1, \ldots, j, \ldots, \ell_{p_j}^{q_j}\}$ the color containing column $\mathsf{C}(:, j)$. As was done earlier, let $\mathsf{C}(:, \boldsymbol{\ell}_{p_j})$ denote the set of columns of C associated with the color $\boldsymbol{\ell}_{p_j}$.

Call $\Gamma_j \equiv \Gamma(\mathsf{B}^T(:, \boldsymbol{\ell}_{p_j}))$ the conflicted index set arising from applying the color $\boldsymbol{\ell}_{p_j}$ to the matrix B^T. The values $\mathsf{C}(i, j)$ and $\hat{\mathsf{C}}(i, j)$ can be computed in pieces related to this Γ_j:

$$\mathsf{C}(i, j) = \sum_{\gamma \in \Gamma_j} \mathsf{A}(i, \gamma)\mathsf{B}^T(\gamma, j) + \sum_{\gamma \notin \Gamma_j} \mathsf{A}(i, \gamma)\mathsf{B}^T(\gamma, j),$$
$$\hat{\mathsf{C}}(i, j) = \sum_{\gamma \in \Gamma_j} \mathsf{A}(i, \gamma)\hat{\mathsf{B}}_c^T(\gamma, j) + \sum_{\gamma \notin \Gamma_j} \mathsf{A}(i, \gamma)\hat{\mathsf{B}}_c^T(\gamma, j).$$

Using the definition of a minimally vacated matrix, we can simplify this second line to

$$\hat{\mathsf{C}}(i,j) = \sum_{\gamma \notin \Gamma_j} \mathsf{A}(i,\gamma)\mathsf{B}^T(\gamma,j),$$

which, combined with the first line, gives

$$\mathsf{C}(i,j) = \sum_{\gamma \in \Gamma_j} \mathsf{A}(i,\gamma)\mathsf{B}^T(\gamma,j) + \hat{\mathsf{C}}(i,j).$$

Since we want $\mathsf{C}(i,j) = \hat{\mathsf{C}}(i,j)$, we must prove that

$$\sum_{\gamma \in \Gamma_j} \mathsf{A}(i,\gamma)\mathsf{B}^T(\gamma,j) = 0, \qquad 1 \le i \le m, \qquad 1 \le j \le n,$$

which we do by proving that

$$\text{for all } \ 1 \le i \le m,\ 1 \le j \le n: \quad \mathsf{A}(i,\gamma) = 0 \ \text{ for each } \ \gamma \in \Gamma_j. \tag{2.4}$$

By the definition of the conflicted index set Γ_j, for every $\gamma \in \Gamma_j$ at least two columns of B^T must have nonzero rows γ. If we choose any two of those and call their indexes s and t, we can write

$$\gamma \in \Gamma_j \ \Rightarrow \ \text{ there exists } \ s,t \in \boldsymbol{\ell}_{p_j},\ s \ne t, \ \text{ such that } \ \mathsf{B}^T(\gamma,s) \ne 0 \text{ and } \mathsf{B}^T(\gamma,t) \ne 0. \tag{2.5}$$

Because $\boldsymbol{\ell}_{p_j}$ is a valid color for C, the set $\mathsf{C}(:,\boldsymbol{\ell}_{p_j})$ is structurally orthogonal:

$$\mathsf{C}(i,s) = 0, \text{ or } \mathsf{C}(i,t) = 0, \qquad 1 \le i \le m.$$

The structural zero $\mathsf{C}(i,s)$ occurs because the vectors $\mathsf{A}(i,:)^T$ and $\mathsf{B}^T(:,s)$ are structurally orthogonal. Applying Lemma 1.2 gives

$$\mathsf{A}(i,\gamma) = 0, \text{ or } \mathsf{B}^T(\gamma,s) = 0, \qquad 1 \le \gamma \le r, \qquad 1 \le i \le m.$$

A similar statement is also true for $\mathsf{C}(i,t)$, so we join these results to state that for every $1 \le \gamma \le r$,

$$\text{either } \ \mathsf{B}^T(\gamma,s) = 0 \text{ or } \mathsf{B}^T(\gamma,t) = 0, \ \text{ or } \ \mathsf{A}(i,\gamma) = 0 \ \text{ for all } 1 \le i \le m. \tag{2.6}$$

The combination of (2.5) and (2.6), along with the knowledge that each column of C appears in only one color, is sufficient to prove (2.4). □

Theorem 2.6 guarantees that the coloring of C is a sufficient tool to perform the sparse-dense matrix product AB^T_{Dense}. While creating the minimally vacated matrix $\hat{\mathsf{B}}^T_c$ is not difficult in practice, most applications do not require it, as shown in Corollary 2.7.

Corollary 2.7. *For the product $\mathsf{C} = \mathsf{AB}^T$, any coloring c of C is a valid coloring of B^T if A has no zero columns.*

Proof. Start with (2.6), which is valid for any two distinct columns $s,t \in \boldsymbol{\ell}_{p_j}$. If A has no zero columns, then $\mathsf{A}(i,\gamma) \ne 0$ for some $1 \le i \le m$, which requires that either $\mathsf{B}^T(\gamma,s) = 0$ or $\mathsf{B}^T(\gamma,t) = 0$ for $1 \le \gamma \le r$ and makes c a valid coloring for B^T. □

Algorithm 3 Computing $\mathsf{C} = \mathsf{AB}^T$ Using Coloring, Basic Version

1: Compute symbolic $\mathsf{C} = \mathsf{AB}^T$; Compute a matrix coloring of C, c
2: Assemble B^T_{Dense} by applying c to B^T
3: Perform the sparse-dense matrix product $\mathsf{C}_{Dense} = \mathsf{AB}^T_{Dense}$
4: Recover C from C_{Dense}

3. Algorithms for Sparse-Sparse Matrix Product Using Coloring. Now that we have laid the foundation, we present Algorithm 3 for computing the product of sparse matrices by an associated sparse-dense product generated through matrix coloring. The remainder of this section analyzes this algorithm and adapt it to address implementation concerns.

One way to analyze Algorithm 3 would be to study the memory traffic. Let n_A and n_B denote the average number of nonzeros per row in A and B, respectively, and let n_{color} be the number of colors in c. From our discussion in Section 1.1, we can surmise that computing AB^T using sparse inner products requires roughly $O((n_\mathsf{A} + n_\mathsf{B})mn)$ memory accesses. On the other hand, AB^T_{Dense} should incur only $O(2n_\mathsf{A}mn_{color})$ memory accesses, which could yield substantial savings if $n_{color} \ll n$. We limit our discussion here to sparse inner products because it offers the most direct comparison, but a comparison of memory traffic with other matrix multiplication techniques will be useful for determining the usefulness of coloring in applications.

Despite the gains from performing a sparse-dense matrix product, costs are incurred by computing the coloring of C, compressing B^T to B^T_{Dense} and then recovering C from C_{Dense}. These eat into the competitive advantage described earlier for the sparse-dense matrix product and must be taken into account when comparing the two algorithms. Rather than try to conduct a "pencil and paper" analysis of this additional complexity, we have performed numerical tests to demonstrate that their cost is not overwhelming. Those results appear in Section 4.

In line 1 of Algorithm 3, we compute the symbolic product of AB^T, which determines the location of the nonzeros in C before their value is computed [19]. Generally, this is used to preallocate space for C, but here we also use this structural description of C to determine a matrix coloring c. We do not discuss the different algorithms for computing matrix colorings here; the effect of two popular choices, which are available in the PETSc library, are compared in Section 4.2.

Line 2 was discussed in Section 2, although here we omit the possible need for a vacated matrix $\hat{\mathsf{B}}^T$ to form B^T_{Dense}. Recall that Corollary 2.7 shows that this vacation is not necessary so long as A has at least one nonzero in each column. The sparse-dense matrix product in line 3 is computed as a sequence of sparse-dense inner products, each of which implements Algorithm 2. The efficient recovery of the sparse matrix C from the computed matrix C_{Dense} is not trivial; it is discussed in Section 3.1.

Lines 2 through 4 contain the so-called numeric portion of the matrix product. For many applications, including the PFLOTRAN example in Section 4.1 and the compactly supported RBF network example in Section 4.3, the nonzero patterns of matrices remain fixed despite varying numeric values; this allows the symbolic component (line 1) to be performed only once while numeric products are computed whenever matrix values are updated. Our efficiency discussions focus on the numeric component.

Although not our immediate focus, performing the sparse-dense product AB^T_{Dense} in lieu of the sparse-sparse product AB^T allows for new optimizations leveraging the

standardized structure of a dense matrix. One optimization that we exploit is the computation of multiple columns of C_{Dense} simultaneously. Each time A is brought into memory, four columns of C_{Dense} are computed, allowing for more flops per call to memory. The number of columns that can be efficiently computed simultaneously is determined by the available memory. Improving the ratio of work per memory access was a key factor in motivating this work, as discussed in Section 1.1, and the topic appears again in Section 4.

3.1. Recovering the Sparse Matrix from the Dense Product. The fourth line of Algorithm 3 decompresses the dense matrix C_{Dense} to the sparse matrix C. Although we have not explicitly stated it previously, more information is needed to perform this decompression than just C_{Dense} and c. In the process of compressing B^T to B^T_{Dense}, multiple columns (previously denoted $\mathsf{B}^T(:, \boldsymbol{\ell}_j)$) are joined to form a single column $\mathsf{B}^T_{Dense}(:, j)$; to undo this process, we must know how to partition $\mathsf{C}_{Dense}(:, j)$ among the columns in the set $\mathsf{C}(:, \boldsymbol{\ell}_j)$.

In practice, during the compression, the coloring is augmented to also store this row data. Referring to the example in Section 2.1, we would augment the coloring

$$c = \{\{1, 4\}, \{2, 5\}, \{3\}, \{6\}\}$$

to include the following row data:

$$c^+ = \{\{1 : \{1, 3\}, 4 : \{4, 5\}\}, \{2 : \{1, 2, 3\}, 5 : \{4\}\}, \{3 : \{1, 3, 4\}\}, \{6 : \{1, 2, 3, 4, 5, 6\}\}\}.$$

Using this augmented coloring c^+, we can decompress a compressed matrix to its sparse form. Note that this is not the only possible form of an augmented coloring; all that is required is sufficient information about the rows of $\mathsf{C}_{Dense}(:, j)$ to recover C. For instance, the row information for the third and fourth colors above could be omitted because only one column is present in that color and thus no ambiguity exists.

The strategy used to populate C from C_{Dense} can contribute significantly to the computational cost of Algorithm 3. A simple approach would be to simply traverse C_{Dense} contiguously and populate C with each nonzero according to c^+. Implementation of this revealed that as the matrix sizes increase, the decompression can consume up to 1/3 of the total execution time; in contrast, the compression of B^T to B^T_{Dense} generally requires a much smaller portion of the total time.

The most direct implementation of the C_{Dense} decompression ignores the fact that values that are very close in C_{Dense} may be in very distant columns of C. This occurs because C_{Dense} is stored in dense *columns* (as is the dense matrix storage standard) but C is a compressed sparse *row* matrix. Unpacking any single column of C_{Dense} may insert nonzero entries throughout C; therefore, decompressing each column of C_{Dense} may cause a traversal through the entire C matrix.

To mitigate this expense, we could fill some block of rows of C all at once, thereby preventing the need for n_{color} passes through C. Results for both decompression techniques are presented in Section 4.1. Algorithm 4 incorporates the changes discussed in this subsection into the coloring-based sparse product algorithm and also notes the potential need for a minimally vacated B^T as described in Theorem 2.6. This algorithm is listed for completeness, to indicate the practical steps that must be taken for implementation; we in general refer instead to Algorithm 3 for simplicity.

3.2. Algorithms for the RAR^T Product. Thus far we have discussed the product (2.1a), but our motivating application is multigrid and it involves the product (2.1b). When Algorithm 3 is adapted for the product $\mathsf{C} = \mathsf{RAR}^T$, two options arise:

Algorithm 4 Computing $\mathsf{C} = \mathsf{AB}^T$ Using Coloring, Practical Version

1: Compute symbolic $\mathsf{C} = \mathsf{AB}^T$
2: Compute c, a matrix coloring of C
3: Vacate B^T using c if needed for compression
4: Assemble B^T_{Dense} by applying c to B^T
5: Augment matrix coloring $c \to c^+$
6: Perform the sparse-dense matrix product $\mathsf{C}_{Dense} = \mathsf{AB}^T_{Dense}$
7: Recover C from C_{Dense} using c^+
8: Populate $m_{blocksize}$ rows of C at once

C can be computed by using two sparse-dense products or one sparse-dense product and one sparse-sparse product. The relative efficiency of these options is affected by the number of colors present in the compression.

The auxiliary matrix W allows us to compute (2.1b) in two steps:

$$\mathsf{W} = \mathsf{AR}^T, \tag{3.1a}$$

$$\mathsf{C} = \mathsf{RW}. \tag{3.1b}$$

We always compute (3.1a) using coloring, but the choice of coloring determines the efficiency of that and subsequent computations. The coloring of W, c_{W}, can be used to implement line 1 in Algorithm 3, at the end of which, a sparse W is returned. Then a sparse-sparse matrix product between R and W is used to compute C. This process is described in Algorithm 5.

Algorithm 5 Computing $\mathsf{C} = \mathsf{RAR}^T$ Using the Coloring of AR^T

1: Use Algorithm 3 to compute $\mathsf{W} = \mathsf{AR}^T$
2: Compute the sparse-sparse CSR matrix product $\mathsf{C} = \mathsf{RW}$

In line 2 of this algorithm, a sparse-sparse matrix product is used instead of the sparse-dense matrix product that we have been developing in this paper. This is counterintuitive because we have focused on replacing sparse-sparse matrix products with coloring-based sparse-dense matrix products, but it may be preferable depending on the number of colors present in C. If C has too many colors, then it may be faster to use a sparse-sparse matrix product to compute it, as discussed in Section 3. Should C have few colors, it may be faster to use Algorithm 6 to compute C.

Algorithm 6 Computing $\mathsf{C} = \mathsf{RAR}^T$ Using the Coloring of C

1: Compute symbolic $\mathsf{C} = \mathsf{RAR}^T$; Compute matrix coloring c_{C}
2: Compress (and vacate, if needed) R^T with c_{C} to form R^T_{Dense}
3: Augment c_{C} to c^+_{C} with the necessary row information
4: Use a sparse-dense matrix product to compute $\mathsf{W}_{Dense} = \mathsf{AR}^T_{Dense}$
5: Use a sparse-dense matrix product to compute $\mathsf{C}_{Dense} = \mathsf{RW}_{Dense}$
6: Recover C from C_{Dense} using c^+_{C}

Algorithm 6 involves two sparse-dense matrix products, each using the coloring c_{C} instead of the coloring c_{W} as was used in Algorithm 5. The cost of this algorithm is tied to the number of colors in C, which can be greater than the number of colors in

W; for many of the multigrid problems we study, C is much more dense than W, and it is less effective to use c_C for these products. Computational results presented in Section 4.1 compare these two algorithms and show that the number of colors present in the coloring is a major factor in the efficiency of the computation.

4. Numerical Experiments. We present four sets of test cases:
- The regional doublet test case from PFLOTRAN [15],
- A three-dimensional linear elasticity PDE test provided in the PETSc library distribution [3],
- Training of a compactly supported RBF network [18], and
- Matrices from the University of Florida Sparse Matrix Collection [7].

The first two tests are chosen because algebraic multigrid is an efficient solver for the linear systems arising in both applications; the third and fourth studies the viability and limitations of the matrix coloring method beyond multigrid applications.

Multigrid is a mathematically optimal method for solving the system of algebraic equations that arise from discretized elliptic boundary value PDEs [21, 22]. We use a geometric-algebraic multigrid method (GAMG) in PETSc, which integrates geometric information into robust algebraic multigrid formulations to yield superior convergence rates of the multigrid solver [2, 3].

GAMG requires the matrix triple products $C = RAR^T$ to be computed on all grid levels in the solver setup phase and the matrix product $C = GG^T$ to be computed for creation of connection graphs; these G matrices have structure derived from that of A and their significance is discussed in [1]. These matrix products, as common computational primitives, constitute a large portion of the entire simulation cost.

The experiments were conducted by using the PETSc library on a Dell Poweredge 1950 server with dual Intel Xeon E5440 quadcore CPUs at 2.83 GHz and 16 GB DDR2-667 memory in 4 channels providing 21 GB/s total memory bandwidth. The machine runs Ubuntu Linux 12.04 64 bit OS. The execution time and floating-point rates were obtained by using one core with the GNU compiler version 4.7.3 and -O optimization. Our performance results were obtained by running the entire test cases and profiling the relevant matrix products. We have done so because standalone benchmarking often produces unreasonably optimistic reports since much of the data is already in cache, which is not the case during an actual simulation.

4.1. Regional Doublet Test Case from PFLOTRAN. We demonstrate the use of matrix coloring on the regional doublet test case [12] from PFLOTRAN, a state-of-the-art code for simulating multiscale, multiphase, multicomponent flow and reactive transport in geologic media. PFLOTRAN solves a coupled system of mass and energy conservation equations for a number of phases, including air, water, supercritical CO_2, and a number of chemical components. PFLOTRAN is built on the PETSc library and makes extensive use of PETSc iterative nonlinear and linear solvers.

The regional doublet test case models variable saturated goundwater flow and solute transport within a hypothetical aquifer measuring 5000 m $\times$ 100 m. We consider only flow problems here because flow solves dominate computation. The governing equation is a system of time-dependant PDEs. PFLOTRAN utilizes finite-volume or mimetic finite-difference spatial discretizations and backward-Euler (fully implicit) timestepping. At each time step, Newton-Krylov methods are used to solve the resulting nonlinear algebraic equations. In all the experiments reported below we have run 35 time steps, which is the minimum needed to resolve the basic physics.

Tables 4.1 and 4.2 show the benefit of using matrix coloring to speed execution time for computing matrix triple products over small to large three-dimensional meshes. Three approaches are compared for computing $\mathsf{C} = \mathsf{RAR}^T$:

- Using no coloring
  - $\mathsf{P}^T\mathsf{AP}$ – stores $\mathsf{P} = \mathsf{R}^T$ in CSR format, and then computes C with sparse *outer* products,
  - RAP – stores $\mathsf{P} = \mathsf{R}^T$ in CSR format, and performs $\mathsf{W} = \mathsf{AP}$ and $\mathsf{C} = \mathsf{RW}$ one row at a time,
  - RAR^T – Algorithm 1 is used to compute each value in C,
- Using the coloring of RAR^T - Algorithm 6, and
- Using the coloring of AR^T - Algorithm 5.

When computing the product of sparse CSR matrices (as occurs in the RAP case above), we use the algorithm [11] that forms a row at a time, i.e., $\mathsf{C}(i,:) = \mathsf{A}(i,:)\mathsf{P}$.

Table 4.1: PFLOTRAN: $\mathsf{C} = \mathsf{RAR}^T$
Fine grid size: $50 \times 25 \times 10$; A: $12,500 \times 12,500$, average nonzeros per row = 7; R: $1,203 \times 12,500$, average nonzeros per row = 34. Times presented are in seconds. The use of coloring in the RAR^T computation is clearly beneficial. Also, the choice of coloring (studying the structure of RAR^T and applying Algorithm 5 or studying AR^T and applying Algorithm 6) plays major role in the algorithm efficiency.

	No Coloring			Coloring	
	$\mathsf{P}^T\mathsf{AP}$	RAP	RAR^T	RAR^T (n_{color}=59)	AR^T (n_{color}=20)
3 Symbolic	.0066	.016	.011	.019	.024
246 Numeric	1.30	1.46	2.39	1.51	.760
Total Time	1.31	1.48	2.40	1.53	.784

Table 4.2: PFLOTRAN: $\mathsf{C} = \mathsf{RAR}^T$
Total computation time (in seconds) presented for increasing fine grid density. The "time_{100}" column implements dense to sparse decompression with 100 row blocks. Algorithm 5 continues to outperform Algorithm 6. Unpacking multiple rows at once during the decompression (as discussed in Section 3.1) speeds the computations for the AR^T coloring but not for the RAR^T coloring.

	No Coloring	Coloring					
	RAR^T	Algorithm 6, using RAR^T			Algorithm 5, using AR^T		
Grid Size	time	n_{color}	time	time_{100}	n_{color}	time	time_{100}
$50 \times 25 \times 10$	2.4	59	1.5	1.6	20	.78	.84
$100 \times 50 \times 20$	28	70	26	27	24	14	12
$200 \times 100 \times 40$	246	84	374	376	25	162	132

The first column of Table 4.1 gives the total number of symbolic and numeric matrix triple products accumulated from all grid levels and all linear iterations during the entire simulation. The symbolic matrix products were computed only during the solver setup phase, while the numeric matrix products were executed for every nonlinear iteration of GAMG. Time spent creating matrix colorings is included in the symbolic row and contributed approximately half the symbolic execution time for the coloring columns. The matrix colorings for this example are created by using

the PETSc default algorithm, largest-first ordering (discussed in Section 4.2). When compared with the repeated execution of numeric matrix products performed during the solve, the time spent on the symbolic products is minimal.

In Table 4.1, we see that the number of colors for matrix RAR^T is 59, whereas AR^T (Column 6) has only 20 colors. This disparity is the result of greater density in C. While both of these are far smaller than 1203, the column size of both R^T and the final sparse matrix product C, the greater sparsity in AR^T demands only a third as many colors and leads to significantly shorter execution time. The last row of Table 4.1 gives the total execution time, that is, the sum of the symbolic and numeric components. It shows that computing $\mathsf{C} = \mathsf{RAR}^T$ by using the matrix coloring of AR^T takes roughly one-third of the execution time of using no coloring.

Table 4.2 presents the same experiments over larger three dimensional meshes; it should be noted that the average number of nonzeros per row for matrices A and R remains unchanged as the grid $50 \times 25 \times 10$, i.e., 7 for A and approximately 34 for R. The use of Algorithm 5 to compute the product continues to outperform standard sparse inner products and Algorithm 6. Here we also consider the dense to sparse decompression in block rows of 100 in addition to conducting the entire decompression in one sweep; this idea was introduced for $\mathsf{C} = \mathsf{AB}^T$ in Section 3.1, and the results are presented in the time_{100} columns. There is a marked benefit when using the block decompression for the AR^T coloring option but no benefit for the RAR^T coloring. This can likely be attributed to the number of colors, since more colors indicates more zero-valued nonzeros during the dense compression and therefore extra work is performed.

Table 4.3: PFLOTRAN: Flop Rate (megaflops/second)
Using coloring greatly increases the flop rate because of the sparse to dense compression. Even though the RAR^T coloring has a higher flop rate, its total computational time (from Table 4.2) is higher because many of the flops performed are unnecessary. The AR^T coloring has few enough colors in it to maintain a high flop rate without excessive unneeded computation involving zero-valued nonzeros. The coloring computation uses the 100 block row decompression.

Fine Meshes	No Coloring	Coloring RAR^T	Coloring AR^T
$50 \times 25 \times 10$	76	1232	724
$100 \times 50 \times 20$	64	936	606
$200 \times 100 \times 40$	59	663	496

Table 4.3 helps clear the computational picture by studying the flop rate for the three methods of computing RAR^T. The floating point computation required for the "No Coloring" option is significantly less than the coloring options because only nonzero matrix entries are involved in the computation; as discussed in Section 1.1, the efficiency (measured by flop rate) suffers as a result.

This issue is alleviated by the coloring because necessary values are stored contiguously in the dense matrix, thereby allowing for more efficient access. That improved efficiency is visible in the two columns of coloring results: both have significantly higher flop rates than the computation without coloring. Following this logic further, the higher flop rate would suggest that the RAR^T coloring is superior to the AR^T coloring. Although computations are happening more quickly there, we know that the total time required for the computation is also greater.

This seeming contradiction is caused by the number of zero-valued nonzeros introduced into the dense matrix. When the sparse matrices are used in the inner product computation, only nonzeros are involved in the computation, but accessing them is a slow proposition. On the other extreme, when the RAR^T coloring is used, too many zeros are involved in the inner products, allowing for more efficient memory access but involving so many zero-valued nonzeros that the total computational time is excessive. The coloring of AR^T has many fewer colors, which reduces the number of superfluous zeros in the dense matrix, and achieves the best overall performance. This is why the flop rate is slightly lower, but the total computational time (shown in Table 4.2) is better.

4.2. Three Dimensional Linear Elasticity PDEs. The tests in this section are found in the PETSc library (see petsc/src/ksp/ksp/examples/tutorials/ex56.c). They model a three-dimensional bilinear quadrilateral (Q1) displacement finite element formulation of linear elasticity PDE defined over a unit cube domain with Dirichlet boundary condition on the side y=0, and load of 1.0 in x + 2y direction on all nodes. Three physical variables are defined at each grid point (i.e., degrees of freedom equals 3). We use the tests to further demonstrate the performance benefit of using matrix colorings and illustrate that the achieved acceleration depends on the matrix nonzero structures and selected matrix coloring.

We apply two matrix coloring algorithms provided as part of MINPACK [17]: Largest-first ordering (LF) [13] and smallest-last ordering (SL) [16] to the grid operator matrices $\mathsf{C} = \mathsf{RAR}^T$ in Table 4.4 and to the matrix product $\mathsf{C} = \mathsf{GG}^T$ (used for connection graphs) in Table 4.5. These tables present the dimensions of the matrices, the execution time spent on computing numeric matrix products, and the number of colors obtained from the matrix colorings. We present only the results using the coloring of AR^T (Algorithm 5) in Table 4.4.

Table 4.4: Elasticity PDE: Execution Time of Numeric $\mathsf{C} = \mathsf{RAR}^T$ (seconds)
As the matrix size increases the number of colors increases more quickly with the LF coloring than the SL coloring. This causes the SL time to scale more effectively, although both colorings outperform the "No Coloring" option.

Matrix Size at Fine Grids (ave. nonzeros/row)		No Coloring	Coloring AR^T			
A	R		LF	n_{color}	SL	n_{color}
3,000 × 3,000 (66)	156 × 3,000 (336)	.022	.0092	48	.0092	48
24,000 × 24,000 (73)	1,122 × 24,000 (456)	.24	.15	66	.15	60
192,000 × 192,000 (77)	8,586 × 192,000 (524)	2.23	1.68	84	1.33	66
648,000 × 648,000 (78)	27,924 × 648,000 (554)	7.86	6.07	90	4.75	69

As the mesh size increases for matrices AR^T, the number of colors produced by the SL algorithm grows slower than that of LF: from 48 to 69 vs. 48 to 90. This results in improved performance using SL colorings as shown in Table 4.4. This stands in contrast to the results for the graph connection matrices GG^T show in Table 4.5 where the LF algorithm generates a more consistent number of colors as the mesh size grows, and outperforms the SL algorithm. These results suggest that no one coloring algorithm is ideal for all circumstances.

4.3. Compactly Supported RBF Networks. Radial basis function (RBF) networks are used in machine learning to predict function values from noisy scattered

Table 4.5: Elasticity PDE: Execution Time of Numeric $\mathsf{C} = \mathsf{G}\mathsf{G}^T$ (seconds)
As the matrix G size grows, the number of colors from the LF coloring remains constant, allowing it to outperform the SL coloring. Again, both coloring choices are faster than the "No Coloring" option.

Matrix Size at Fine Grids	No Coloring	Coloring			
G (ave. nonzeros/row)		LF	n_{color}	SL	n_{color}
1,000 × 1,000 (20)	.012	.0026	125	.0029	136
8,000 × 8,000 (23)	.13	.038	125	.043	149
64,000 × 64,000 (25)	1.22	.42	125	.51	158
216,000 × 216,000 (26)	4.31	1.49	125	1.81	161

data [18]; the network is a linear combination of RBFs which is trained by solving an optimization problem. To solve that problem, which balances fidelity of the network with its regularity (controlled by $\mu \in \mathbb{R}$), the linear system

$$(\mathsf{K}\mathsf{K}^T + \mu \mathsf{I}_m)\boldsymbol{c} = \mathsf{K}\boldsymbol{y} \tag{4.1}$$

is solved, where $\mathsf{K} \in \mathbb{R}^{m \times n}$ is populated by RBF values centered at m nodes and evaluated at the n data locations. The regularity parameter μ prevents oscillations caused by noise, and $\boldsymbol{y} \in \mathbb{R}^n$ are the function values evaluated at the n data locations.

Many common RBFs (e.g., Gaussians) will produce dense K matrices, but compactly supported RBFs such as the Wendland or Wu functions will produce sparse K matrices [23]. As a result, the $\mathsf{K}\mathsf{K}^T$ product will be a sparse matrix product. Furthermore, compactly supported RBFs are often of the form

$$K(\boldsymbol{x}, \boldsymbol{z}; \ell, \varepsilon) = \phi_\ell(\varepsilon r)(1 - \varepsilon r)_+, \qquad r = \|\boldsymbol{x} - \boldsymbol{z}\|_2, \tag{4.2}$$

where ℓ is a *smoothness parameter* such that ϕ_ℓ is a polynomial of degree ℓ and ε is a *shape parameter* which allows for manipulation of the support of the RBF K. Because the nonzero structure of the matrix is determined by the ε term, different smoothness parameters can by tested without changing the structure. This makes compactly supported RBF networks an application which could benefit from matrix multiplication through coloring.

Fig. 4.1 displays the results of computing $\mathsf{K}\mathsf{K}^T$ for a sample of $n = 200000$ Halton points in 2D [8] with variable m Halton kernel centers (the number of rows of K) and variable ε (larger ε yields a sparser matrix). A clear middle region of ε and m values exists where the **Coloring** computation is more efficient than a **Direct** computation involving sparse inner products and the explicit storage of the **Transpose** for use in a CSR matrix product. Matrices associated with the bottom-right of this graph are small and sparse and movement towards the top-left yields larger and denser matrices; the coloring algorithm is fastest in the transition region.

4.4. Matrices from the University of Florida Sparse Matrix Collection. To examine the efficiency of applying matrix coloring to the matrices beyond the multigrid applications, we selected matrices from the University of Florida Sparse Matrix Collection [7]:

1. `SNAP/amazon0302` - commerce network: A is 262111 × 262111 with 1234877 nonzeros.
2. `Mittelmann/watson_2` - linear programming: A is 352013 × 677224 with 1846391 nonzeros.

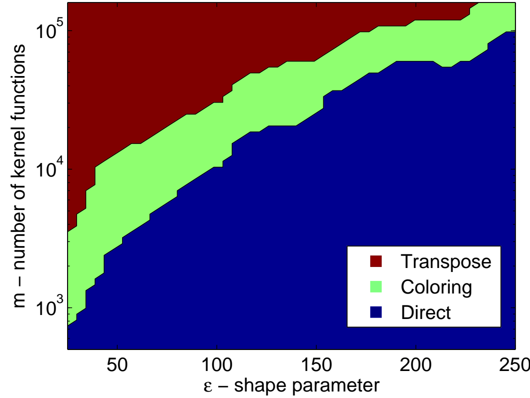

Fig. 4.1: These experiments compare the cost of the numerical matrix product $\mathsf{K}\mathsf{K}^T$ required for (4.1) as computed with three different methods: the **Coloring** method using Algorithm 4, a **Direct** sparse inner product computation involving Algorithm 1 (preferable in the lower-right region) and an explicit computation and storage of the **Transpose** followed by the product of two CSR matrices (preferable in the upper-left region). For a variety of RBF matrix widths (x axis) and shape parameters (y axis), there exists a middle region where the coloring computation is preferable.

3. `JGD_Homology/ch7-8-b5` - combinatorial: A is 141120×141120 with 846720 nonzeros.
4. `SNAP/wiki-Vote` - social network: A is 8297×8297 with 103689 nonzeros.
5. `SNAP/roadNet-CA` - transportation: A is 1971281×1971281 with 5533214 nonzeros.
6. `GHS_psdef/bmwcra_1` - structural: A is 148770×148770 with 10641602 nonzeros.
7. `Gleich/usroads` - transportation: A is 129164×129164 with 330870 nonzeros.
8. `DIMACS10/144` - undirected graph: A is 144649×144649 with 2148786 nonzeros.
9. `DNVS/m_t 1` - structural: A is 97578×97578 with 9753570 nonzeros.
10. `Norris/torso2` - 2D/3D: A is 115967×115967 with 1033473 nonzeros.
11. `Williams/cant` - finite elements: A is 62451×62451 with 4007383 nonzeros.
12. `GHS_psdef/s3dkq4m2` - structural: A is 90449×90449 with 4427725 nonzeros.

For these matrices $\mathsf{A} \in \mathbb{R}^{m \times n}$, we compute $\mathsf{C} = \mathsf{A}\mathsf{A}^T \in \mathbb{R}^{m \times m}$. To effectively study the benefit of computing sparse inner products with this coloring approach, we introduce a ratio which describes how many additional nonzeros are introduced into C_{Dense}:

$$\text{compression ratio} \;=\; \frac{\#\text{ nonzeros of } \mathsf{C}_{Dense}}{\#\text{ nonzeros of } \mathsf{C}} \tag{4.3}$$

Because C_{Dense} is stored as a dense matrix, # nonzeros of $\mathsf{C}_{Dense} = m\, n_{color}$. This ratio appears in Table 4.6; if the dense matrix has introduced few "zero-valued nonzeros" then the ratio tends towards 1.

Table 4.6 suggests that products of the form $\mathsf{C} = \mathsf{A}\mathsf{A}^T$ benefit from using the coloring strategy rather than sparse inner products primarily if the ratio in (4.3) is sufficiently low. In general, if this ratio is too high, then too much additional work

Table 4.6: For the A matrices listed above, $\mathsf{C} = \mathsf{AA}^T$ is computed, and the numerical computation time (but not the symbolic time) is listed in seconds. When the coloring strategy introduces few nonzeros, its cost may be less than the sparse inner product computation in the "No Coloring" column. This penalty for using a dense matrix is tracked in the "ratio" column, where a value of 1 would indicate no nonzeros are introduced.

Matrix		No Coloring	Coloring			
name	avg. nnz per row	time	time	n_{color}	m/n_{color}	ratio
`SNAP/amazon0302`	5	1.63	25.24	2032	129	61.72
`Mittelmann/watson_2`	5	0.44	1.82	228	1544	15.14
`JGD_Homology/ch7-8-b5`	6	2.74	8.34	1745	81	12.24
`SNAP/wiki-Vote`	12	2.03	1.68	4819	2	8.47
`SNAP/roadNet-CA`	3	0.64	0.94	19	103751	3.64
`GHS_psdef/bmwcra_1`	72	13.01	14.26	1155	129	3.46
`Gleich/usroads`	2.5	.035	.049	16	8073	3.44
`DIMACS10/144`	15	1.90	1.50	191	757	3.21
`DNVS/m_t1`	100	12.64	10.17	945	103	2.52
`Norris/torso2`	9	0.16	0.14	42	2761	1.70
`Williams/cant`	64	6.09	2.07	401	155	1.44
`GHS_psdef/s3dkq4m2`	50	2.66	1.22	178	508	1.21

is created by working with a dense matrix to take advantage of the better flop rate for sparse-dense inner products discussed in Table 4.3; note, however, that this result is not always true, as the `SNAP/roadNet-PA` matrix is better compressed than the `SNAP/wiki-Vote`, but performs much worse. The average number of nonzeros does not seem immediately useful as a tool to predict of the speed of this coloring strategy, in contrast to our initial beliefs. Similarly, the average number of columns per color, measured by m/n_{color}, seems to be an unhelpful diagnostic tool, suggesting that developing a strategy to decide the viability of coloring is nontrivial.

We note that matrices with origins in network science (including `Gleich/usroads` and `SNAP/amazon0302`) seem to rarely benefit from the coloring strategy. This may be caused by the structure of such matrices, e.g., power law or small network graphs. Possible remedies for this issue include working with A as a block matrix to exploit colorings within blocks, or an alternate coloring strategy specifically for such applications.

Another shortcoming of this coloring strategy that must be explicitly mentioned is the additional memory cost of using it on matrices with a high "ratio" value. When this value is greater than two, more memory must be allocated for the compressed matrix than was required for the original matrix, which is a significant drawback that can also cripple the efficiency of this method.

5. Conclusions and Future Work. Earlier research has accelerated the computation of sparse Jacobians by using matrix coloring to evaluate several columns simultaneously. This work has been adapted to accelerate sparse matrix products computed with sparse inner products by applying a matrix coloring to instead compute a related compressed sparse-dense matrix product. We have proved that for the product $\mathsf{C} = \mathsf{AB}^T$, the matrix coloring of C is always a viable choice for compressing B^T, although slight modifications may be necessary if A has any zero columns.

Algorithms for both (2.1a) and (2.1b) were proposed, as well as considerations for practical implementation. Numerical results suggested that simulations using multigrid that computed sparse matrix products through inner products can be accelerated significantly through the use of coloring. Those results can be improved by fine tuning the data traffic during the decompression of the computed dense matrix to the desired sparse matrix. We also studied the effect of the choice of coloring algorithm on the efficiency of the compression and found that different algorithms perform better in different circumstances.

Future adaption of this coloring approach to parallel matrices will inherit much of the theory, but the added cost of data traffic may require some retooling or reorganization. Additionally, it would be valuable to study the use of this method on block matrices, a common structure for many applications. Beyond multigrid, our motivating application, we believe that applications from the discrete math community (e.g., breadth-first search) will benefit from our coloring algorithm.

The transition to a sparse-dense matrix product opens the door to numerous optimizations that are unavailable in the sparse-sparse setting. For instance, our current software computes four columns of the dense matrix each time the sparse matrix is loaded into memory, thereby reducing the total cost of accessing the sparse matrix. Another improvement we hope to consider is the interlacing of several columns of the dense matrix during the product to allow for more computation per sparse matrix cache miss. Decompressing the resulting product requires new techniques, so more work is required to take advantage of this and other previously unavailable opportunities.

We plan to incorporate the block dense to sparse decompression discussed in Section 3.1 into the PETSc library for faster computation of finite-difference Jacobians using coloring. Another interesting advance might be to develop a new coloring algorithm that incorporates the cost of decompression when deciding how to organize colors. A coloring algorithm that accelerates the decompression would benefit both the matrix multiplication setting and the evaluation of finite difference Jacobians. Furthermore, the value of this algorithm lies primarily in situations where the coloring can be reused for many products with the same sparsity structure, and thus its practicality is a function of the ratio of the cost of finding the coloring to the cost of the numerical computation; as a consequence, this algorithm would not be appropriate for a single RAR^T product. If an algorithm for computing the coloring were designed to take advantage of the AB^T structure we have studied here, it would improve the viability of this approach by decreasing the setup cost.

Acknowledgments. The authors were supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357. The authors sincerely thank the referees for their constructive comments, which strengthened the presentation of this manuscript. We extend our thanks to Jed Brown, Lois Curfman McInnes and Charles Van Loan for their help and support, and Glenn Hammond for providing us the PFLOTRAN test case.

REFERENCES

[1] M. F. Adams, *Algebraic multrigrid methods for constrained linear systems with applications to contact problems in solid mechanics*, Numerical Linear Algebra with Applications, 11 (2004), pp. 141–153.

[2] M. F. ADAMS, H. H. BAYRAKTAR, T. M. KEAVENY, AND P. PAPADOPOULOS, *Ultrascalable implicit finite element analyses in solid mechanics with over a half a billion degrees of freedom*, in ACM/IEEE Proceedings of SC2004: High Performance Networking and Computing, 2004. Gordon Bell Award.
[3] S. BALAY, J. BROWN, K. BUSCHELMAN, V. EIJKHOUT, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNES, B. F. SMITH, AND H. ZHANG, *PETSc users manual*, Tech. Report ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013.
[4] A. BULUC AND J. R. GILBERT, *Parallel sparse matrix-matrix multiplication and indexing: Implementation and experiments*, SIAM Journal on Scientific Computing, (2012).
[5] A. R. CURTIS, M. J. D. POWELL, AND J. K. REID, *On the estimation of sparse Jacobian matrices*, J. Inst. Maths Applics, (1974), pp. 117–119.
[6] S. DALTON, N. BELL, AND L. N. OLSON, *Optimizing sparse matrix-matrix multiplication for the GPU*, tech. report, University of Illinois Urbana-Champaign, 2014.
[7] T. A. DAVIS AND Y. HU, *The University of Florida sparse matrix collection*, ACM Transactions on Mathematical Software, 38 (2011), pp. 1–25.
[8] G. E. FASSHAUER, *Meshfree Approximation Methods with* MATLAB, vol. 6 of Interdisciplinary Mathematical Sciences, World Scientific Publishing Co., Singapore, 2007.
[9] A. H. GEBREMEDHIN, F. MANNE, AND A. POTHEN, *What color is your Jacobian? Graph coloring for computing derivatives*, SIAM REVIEW, 47 (2005), pp. 629–705.
[10] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations (4th ed.)*, Johns Hopkins University Press, Baltimore, MD, 2012.
[11] F. G. GUSTAVSON, *Two fast algorithms for sparse matrices: multiplication and permuted transposition*, ACM Transactions on Mathematical Software, 4 (1978), pp. 250–269.
[12] G. E. HAMMOND, P. C. LICHTNER, AND R. T. MILLS, *Evaluating the performance of parallel subsurface simulators: An illustrative example with PFLOTRAN*, Water Resources Research, (2013).
[13] W. KLOTZ, *Graph coloring algorithms*, Mathematics Report, (2002), pp. 1–9.
[14] M. KUBALE, *Graph Colorings*, Contemporary Mathematics (American Mathematical Society) v. 352, American Mathematical Society, 2004.
[15] P. LICHTNER ET AL., *PFLOTRAN project*. http://ees.lanl.gov/pflotran/.
[16] D. W. MATULA AND L. L. BECK, *Smallest-last ordering and clustering and graph coloring algorithms*, J. ACM, 30 (1983), pp. 417–427.
[17] J. J. MORÉ, D. C. SORENSON, B. S. GARBOW, AND K. E. HILLSTROM, *The MINPACK project*, in Sources and Development of Mathematical Software, W. R. Cowell, ed., 1984, pp. 88–111.
[18] M. J. L. ORR, *Introduction to radial basis function networks*, tech. report, University of Edinburgh, Centre for Cognitive Sciences, 1996.
[19] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, PWS Publishing Company, Boston, MA, 1996.
[20] B. SMITH AND H. ZHANG, *Sparse triangular solves for ILU revisited: Data layout crucial to better performance*, International J. High Performance Computing Applications, 25 (2011), pp. 386–391.
[21] K. STÜBEN, *A review of algebraic multigrid*, J. Comput. Appl. Math., 128 (2001), pp. 281–309.
[22] U. TROTTENBERG, C. W. OOSTERLEE, AND A. SCHULLER, *Multigrid*, Elsevier Science, 2000.
[23] H. WENDLAND, *Scattered Data Approximation*, vol. 17 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, 2005.

Appendix. *Proof of Theorem 1.6.* Each column of C_{Dense} is created from a single color, and no column of C appears in more than one color. Thus, proving that any column of C_{Dense} is unique proves that C is unique. We will prove that each value of the kth column of C_{Dense} is unique.

If the value $\mathsf{C}_{Dense}(j,k) = 0$, then $\sum_{i=1}^{q_k} |\mathsf{C}(j, \ell_k^i)| = 0$, which could occur only if $\mathsf{C}(j, \ell_k^1) = \ldots = \mathsf{C}(j, \ell_k^{q_k}) = 0$. This would prevent having a conflicting nonzero value for that location, so all zero values in the kth column are unique.

If the value $\mathsf{C}_{Dense}(j,k) \neq 0$, then at least one value in $\{\mathsf{C}(j, \ell_k^1), \ldots, \mathsf{C}(j, \ell_k^{q_k})\}$ is nonzero. If exactly one value is nonzero, then the value $\mathsf{C}_{Dense}(j,k)$ is unique. We must prove that only one of these q_k values is nonzero; we begin by assigning the nonzero value to the ν_k index, namely, $\mathsf{C}_{Dense}(j, \ell_k^{\nu_k}) \neq 0$.

Because c is a valid matrix coloring of C, the columns $\{\mathsf{C}(:, \ell_k^1), \ldots, \mathsf{C}(:, \ell_k^{q_k})\}$ must

be structurally orthogonal. Lemma 1.2 tells us that for the jth row,

$$\mathsf{C}(j, \ell_k^s)\mathsf{C}(j, \ell_k^{\nu_k}) = 0, \qquad 1 \leq s \leq q_k, \quad s \neq \nu_k.$$

Since we have assumed that $\mathsf{C}_{Dense}(j, \ell_k^{\nu_k}) \neq 0$, this lemma demands that $\mathsf{C}(j, \ell_k^s) = 0$ for $1 \leq s \leq q_k$, $s \neq \nu_k$. Therefore, every value in the kth column of C_{Dense} is uniquely determined. □

Proof of Lemma 2.3. If $\Gamma = \emptyset$, then the vectors in $\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_q\}$ are already structurally orthogonal. Otherwise, we must prove that

$$|\hat{\boldsymbol{u}}_i|^T |\hat{\boldsymbol{u}}_j| = 0, \qquad 1 \leq i, j \leq q, \qquad i \neq j.$$

Let us simplify the notation for this proof by defining $\Gamma \equiv \Gamma(\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_q\})$. For any $i \neq j$, the inner product can be separated into two components,

$$\begin{aligned} |\hat{\boldsymbol{u}}_i|^T |\hat{\boldsymbol{u}}_j| &= \sum_{\gamma \in \Gamma} |\hat{\boldsymbol{u}}_i(\gamma)||\hat{\boldsymbol{u}}_j(\gamma)| + \sum_{\gamma \notin \Gamma} |\hat{\boldsymbol{u}}_i(\gamma)||\hat{\boldsymbol{u}}_j(\gamma)| \\ &= \sum_{\gamma \in \Gamma} |0||0| + \sum_{\gamma \notin \Gamma} |\boldsymbol{u}_i(\gamma)||\boldsymbol{u}_j(\gamma)| = \sum_{\gamma \notin \Gamma} |\boldsymbol{u}_i(\gamma)||\boldsymbol{u}_j(\gamma)|. \end{aligned}$$

Because the conflicted index set includes indices such that $\boldsymbol{u}_i(\gamma)\boldsymbol{u}_j(\gamma) \neq 0$, any $\gamma \notin \Gamma$ must have $\boldsymbol{u}_i(\gamma)\boldsymbol{u}_j(\gamma) = 0$, leaving the summation above equal to zero. □

Government License. The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.