EXPOSING FINE-GRAINED PARALLELISM IN ALGEBRAIC MULTIGRID METHODS

NATHAN BELL[†], STEVEN DALTON[‡], AND LUKE N. OLSON[§]

Abstract. Algebraic multigrid methods for large, sparse linear systems are a necessity in many computational simulations, yet parallel algorithms for such solvers are generally decomposed into coarse-grained tasks suitable for distributed computers with traditional processing cores. However, accelerating multigrid on massively parallel throughput-oriented processors, such as the GPU, demands algorithms with abundant *fine-grained* parallelism. In this paper, we develop a parallel algebraic multigrid method which exposes substantial fine-grained parallelism in both the construction of the multigrid hierarchy as well as the cycling or solve stage. Our algorithms are expressed in terms of scalable parallel primitives that are efficiently implemented on the GPU. The resulting solver achieves an average speedup of over 2× in the setup phase and around 6× in the cycling phase when compared to a representative CPU implementation.

Key words. algebraic multigrid, parallel, sparse, gpu, iterative

AMS subject classifications. 65-04, 68-04, 65F08, 65F50, 68W10

1. Introduction.

Throughput-oriented processors, such as graphics processing units (GPUs), are becoming an integral part of many high-performance computing systems. In contrast to traditional CPU architectures, which are optimized for completing scalar tasks with minimal latency, modern GPUs are tailored for parallel workloads that emphasize total task throughput [17]. Therefore, harnessing the computational resources of the such processors requires programmers to decompose algorithms into thousands or tens of thousands of separate, fine-grained threads of execution. Unfortunately, the parallelism exposed by previous approaches to algebraic multigrid is too coarse-grained for direct implementation on GPUs.

Algebraic multigrid methods solve large, sparse linear systems $Ax = b$ by constructing a hierarchy of grid levels directly from the matrix A. In this paper, we consider effective implementations of algebraic multigrid methods for GPUs. We study the components that comprise the two distinct phases in AMG (i.e., the setup and solve phases) and demonstrate how they can be decomposed into scalable parallel primitives.

Parallel approaches to multigrid are plentiful. Algebraic multigrid methods have been successfully parallelized on distributed-memory CPU clusters using MPI [12, 10] and more recently with a combination of MPI and OpenMP [2], to better utilize multi-core CPU nodes. While such techniques have demonstrated scalability to large numbers of processors, they are not immediately applicable to the GPU. In particular, effective use of GPUs requires substantial *fine-grained* parallelism at all stages of the computation. In contrast, the parallelism exposed by existing methods for distributed-memory clusters of traditional cores is comparably coarse-grained and cannot be scaled down to arbitrarily small subdomains. Indeed, coarse-grained parallelization strategies are qualitatively different than fine-grained strategies.

For example, it is possible to construct a successful parallel coarse-grid selection algorithm by partitioning a large problem into sub-domains and applying an effective,

[†]NVIDIA Research, nbell@nvidia.com, http://www.wnbell.com

[‡]Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, dalton6@illinois.edu

[§]Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801, lukeo@illinois.edu, http://www.cs.illinois.edu/homes/lukeo

serial heuristic to select coarse-grid nodes on the the interior of each sub-domain, followed by a less-effective but parallel heuristic to the interfaces between sub-domains [23]. An implicit assumption in this strategy is that the interiors of the partitions (collectively) contain the vast majority of the entire domain, otherwise the serial heuristic has little impact on the output. Although this method can be scaled down to arbitrarily fine-grained parallelism in principle, the result is qualitatively different. In contrast, the methods we develop do not rely on partitioning and expose parallelism to the finest granularity — i.e., one thread per matrix row or one thread per nonzero entry.

Geometric multigrid methods were the first to be parallelized on GPUs [20, 9, 34]. These "GPGPU" approaches, which preceded the introduction of the CUDA and OpenCL programming interfaces, programmed the GPU through existing graphics application programming interfaces (APIs) such as OpenGL and Direct3d. Subsequent works demonstrated GPU-accelerated geometric multigrid for image manipulation [26] and CFD [13] problems. Previous works have implemented the cycling stage of algebraic multigrid on GPUs [19, 22], however hierarchy construction remained on the CPU. A parallel aggregation scheme is described in [35] that is similar to ours based on maximal independent sets, while in [1] the effectiveness of parallel smoothers based on sparse matrix-vector products is demonstrated. Although these works were implemented for distributed CPU clusters, they are amenable to fine-grained parallelism as well.

In Section 1.2, we review the components of the setup and solve phases of AMG, noting the principal challenge in targeting GPU acceleration. Our approach to the GPU is to describe the components of multigrid in terms of parallel primitives, which we define in Section 2. In Section 3 we detail our specific approach to exposing fine-grained parallelism in the components of the setup phase, and in Section 4, we highlight the value of the sparse matrix-vector product in the computation. In Section 5 we discuss several performance results of our method on the GPU in comparison to the efficiency we observe on a standard CPU. Finally, in the Appendix we provide additional details of a parallel aggregation method.

1.1. Background. Multigrid methods precondition large, sparse linear systems of equations and in recent years have become a robust approach for a wide range of problems. One reason for this increase in utility is the trend toward more *algebraic* approaches. In the classical, *geometric* form of multigrid, the performance relies largely on specialized smoothers, and a hierarchy of grids and interpolation operators that are predefined through the geometry and physics of the problem. In contrast, algebraic-based multigrid methods (AMG) attempt to automatically construct a hierarchy of grids and intergrid transfer operators without explicit knowledge of the underlying problem — i.e., directly from the linear system of equations [32, 37]. Removing the requirement of geometric information increases the applicability of AMG to problems with complicated domains and unstructured grids, but also places increased demands on the sparse matrix algorithms.

In the remainder of this section, we outline the basic components of AMG in an *aggregation* context [37] and highlight the necessary sparse matrix computations used in the process. We restrict our attention to aggregation methods because of the flexibility in the construction, however our development also extends to classical AMG methods based on coarse-fine splittings [32].

1.2. Components of Algebraic Multigrid. Central to algebraic-based multigrid methods is the concept of *algebraically* smooth error. That is, error modes *not*

sufficiently reduced by a relaxation method such as weighted Jacobi, Gauss-Seidel, Chebyshev, or Kaczmarz, are considered algebraically smooth and must be handled by coarse-grid correction. Aggregation-based methods are designed to accurately represent such low-energy modes *by construction*. Specifically, the interpolation operators, which transfer solutions between the coarse and find grid, are defined by insisting that a given set of low-energy modes on the fine grid, referred to as *near-nullspace candidates*, are represented exactly on the coarse grid. The performance of AMG relies on a compatible collection of relaxation operators, coarse grid operators, and interpolation operators as well as the efficient construction of these operations. In this section we outline the components of aggregation-based AMG that we consider for construction on the GPU.

Aggregation-based AMG requires a *a priori* knowledge or prediction of the near-nullspace that represent the low-energy error. For an $n \times n$ symmetric, positive-definite matrix problem $Ax = b$, these m modes are denoted by the $n \times m$ column matrix B. Generally, the number of near-nullspace modes (m) is a small, problem-dependent constant. For example, the scalar Poisson problem requires only a single near-nullspace mode while 6 rigid body modes are needed to solve three-dimensional elasticity problems. We also denote the $n \times n$ problem as the *fine* level and label the indices $\Omega_0 = \{0, \ldots, n\}$ as the fine *grid*. From A, b, and B, the components of the solver are defined through a setup phase, and include grids Ω_k, interpolation operators P_k, restriction operators R_k, relaxation error propagation operators, and coarse representations of the matrix operator A_k, all for each level k. We denote index M as the maximum level — e.g., $M = 1$ is a two-grid method.

1.2.1. Setup Phase. We follow a setup phase that is outlined in Algorithm 1. The following sections detail each of the successive routines in the setup phase: `strength`, `aggregate`, `tentative`, `prolongator`, and the triple matrix Galerkin product. One of the goals of this paper is to systematically consider the sparse matrix operations in Lines 1-6.

Algorithm 1: AMG Setup: `setup`

parameters: A, sparse problem matrix
B, m low-energy vectors

return: $A_0, \ldots, A_M$, hierarchy of matrix operators
$P_0, \ldots, P_{M-1}$, hierarchy of interpolation matrices

$A_0 \leftarrow A$, $B_0 \leftarrow B$
for $k = 0, \ldots, M$

1 $C_k \leftarrow \texttt{strength}(A_k)$ {strength-of-connection}
2 $Agg_k \leftarrow \texttt{aggregate}(C_k)$ {construct coarse aggregates}
3 $T_k,\ B_{k+1} \leftarrow \texttt{tentative}(Agg_k, B_k)$ {form tentative interpolation}
4 $P_k \leftarrow \texttt{prolongator}(A_k, T_k)$ {improve interpolation}
5 $R_k \leftarrow P_k^T$ {transpose}
6 $A_{k+1} \leftarrow R_k A_k P_k$ {coarse matrix, triple-matrix product}

1.2.2. Strength-of-connection. In order to construct coarse grids that capture the most important features of the fine grid it is necessary to define the concept of

the relative strength-of-connection between points in the fine-level matrix graph. A vertex i in the fine matrix graph that strongly influences or strongly depends on a neighboring vertex j typically has a large, relative edge weight. As a result, the traditional approach for aggregation schemes is to identify two points i and j as strongly connected if they satisfy

$$|A(i,j)| > \theta\sqrt{|A(i,i)A(j,j)|}. \tag{1.1}$$

This concise statement yields a connectivity graph represented by sparse matrix C_k. Algorithm 2 describes the complete strength-of-connection algorithm for the COO matrix format. A parallel implementation of this algorithm is discussed in Section 3.1.

Algorithm 2: Strength of connection: `strength`

parameters: $A_k \equiv (I, J, V)$, COO sparse matrix
return: $C_k \equiv (\hat{I}, \hat{J}, \hat{V})$, COO sparse matrix

$\mathcal{M} = \{0, \ldots, nnz(A)\}$
$D \leftarrow 0$
for $n \in \mathcal{M}$ {extract diagonal}
  if $I_n = J_n$
    $D(I_n) \leftarrow V_n$
for $n \in \mathcal{M}$ {check strength}
  if $|V_n| > \theta\sqrt{|D(I_n)| \cdot |D(J_n)|}$
    $(\hat{I}_{\hat{n}}, \hat{J}_{\hat{n}}, \hat{V}_{\hat{n}}) \leftarrow (I_n, J_n, V_n)$

1.2.3. Aggregation. An aggregate or grouping of nodes is defined by a root node i and its neighborhood — i.e., all points j, for which $C(i,j) \neq 0$, where C is a strength matrix. The standard aggregation procedure consists of two phases:

1. For each node i, if i and each of its strongly connected neighbors are not yet aggregated, then form a new aggregate consisting of i and its neighbors.
2. For each remaining unaggregated node i, sweep i into an adjacent aggregate.

The first phase of the algorithm visits each node and attempts to create disjoint aggregates from the node and its 1-ring neighbors. It is important to note that first phase is a greedy approach and therefore sensitive to the order in which the nodes are visited. We revisit this artifact in Section 3.2, where we devise a parallel aggregation scheme that mimics the standard sequential algorithm up to a reordering of the nodes.

Nodes that are not aggregated in the first phase are incorporated into an adjacent aggregate in the second phase. By definition, each unaggregated node must have at least one aggregated neighbor (otherwise it could be the root of a new aggregate) so all nodes are aggregated after the second phase. When an unaggregated node is adjacent to two or more existing aggregates, an arbitrary choice is made. Alternatively, an aggregate with the largest/smallest index or the aggregate with the fewest members, etc., could be selected. Figure 1.1 illustrates a typical aggregation pattern for structured and unstructured meshes.

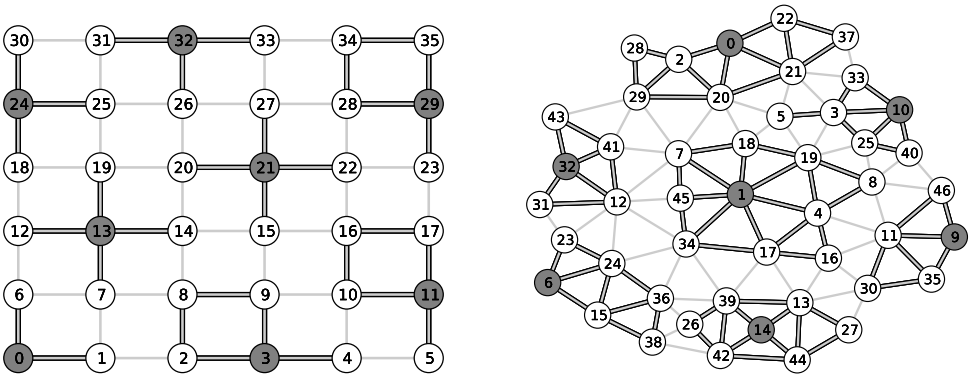

(a) Structured Mesh Aggregates (b) Unstructured Mesh Aggregates

FIG. 1.1. *Example of a mesh (gray) and aggregates (outlined in black). Nodes are labeled with the order in which they are visited by the sequential aggregation algorithm and the root nodes, selected in the first phase of the algorithm, are colored in gray. Nodes that are adjacent to a root node, such as nodes 1 and 6 in 1.1a are aggregated in phase 1. Nodes that are not adjacent to a root node, such as nodes 8, 16, and 34 in 1.1a are aggregated in second phase.*

The aggregation process results in a sparse matrix Agg which encodes the aggregates using the following scheme,

$$Agg(i,j) = \begin{cases} 1 & \text{if the } i^{th} \text{ node is contained in the } j^{th} \text{ aggregate} \\ 0 & \text{otherwise.} \end{cases} \tag{1.2}$$

For example, the aggregation depicted in Figure 1.2, which associates nodes i_1 and i_2 with the first aggregate and nodes i_3, i_4 and i_5 with the second aggregate, would be encoded as,

$$Agg = \begin{array}{c} \\ i_1 \\ i_2 \\ i_3 \\ i_4 \\ i_5 \end{array} \begin{array}{c} \begin{array}{cc} j_1 & j_2 \end{array} \\ \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{bmatrix} \end{array}. \tag{1.3}$$

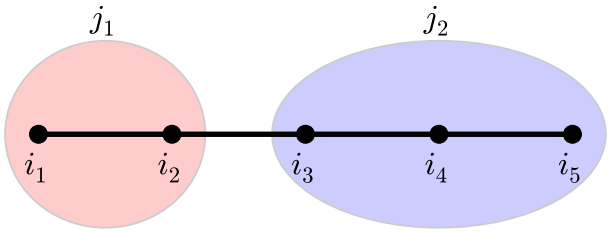

FIG. 1.2. *Nodes of a graph divided into aggregates* $\{i_1, i_2\}$ *and* $\{i_3, i_4, i_5\}$.

1.2.4. Construction of the Tentative Prolongator. From an aggregation, represented by Agg, and a set of coarse-grid candidate vectors, represented by B, a

tentative interpolation operator is defined. The tentative interpolation operator or prolongator matrix T is constructed such that each row corresponds to a grid point and each column corresponds to an aggregate. When there is only one candidate vector, the sparsity pattern of the tentative prolongator is exactly the same as Agg. For example,

$$T = \begin{bmatrix} * & \\ * & \\ & * \\ & * \\ & * \end{bmatrix}, \tag{1.4}$$

is the sparsity pattern for the aggregation illustrated in Figure 1.2. With the sparsity pattern induced by Agg_k, the numerical entries of T_k are defined by the conditions,

$$B_k = T_k B_{k+1}, \qquad T_k^T T_k = I, \tag{1.5}$$

which imply that (1) the near-nullspace candidates lie in the range of T_k and (2) that columns of T_k are orthonormal. For example, the matrices

$$B_k = \begin{bmatrix} B_{1,1} \\ B_{2,1} \\ B_{3,1} \\ B_{4,1} \\ B_{5,1} \end{bmatrix}, \qquad T_k = \begin{bmatrix} B_{1,1}/C_1 & \\ B_{2,1}/C_1 & \\ & B_{3,1}/C_2 \\ & B_{4,1}/C_2 \\ & B_{5,1}/C_2 \end{bmatrix}, \qquad B_{k+1} = \begin{bmatrix} C_1 \\ C_2 \end{bmatrix}, \tag{1.6}$$

satisfy the interpolation ($B_k = T_k B_{k+1}$) and orthonormality conditions ($T_k^T T_k = I$), using the scaling factors $C_1 = ||[B_{1,1}, B_{2,1}]||$ and $C_2 = ||[B_{3,1}, B_{4,1}, B_{5,1}]||$.

Although we only consider the case of a single candidate vector in this paper, for completeness we note that when B_k contains $m > 1$ low-energy vectors, the tentative prolongator takes on a block structure. For example, with $m = 2$ each nonzero in Agg is expanded into a 2-by-1 dense block,

$$T = \begin{bmatrix} * & * & \\ * & * & \\ & & * & * \\ & & * & * \\ & & * & * \end{bmatrix}. \tag{1.7}$$

An important component of the setup phase is to express these operations efficiently on the GPU, which we detail in Section 3.4.

1.2.5. Prolongator Smoothing. The tentative prolongation operator is a direct attempt to enforce the range of interpolation to coincide with the (user-provided) near null-space modes. This has limitations however, since the modes may not accurately represent the "true" near null-space and since the interpolation is still only local, and thus limited in accuracy. One approach to improving the properties of interpolation is to smooth the columns of the tentative prolongation operator. With weighted-Jacobi smoothing, for example, the operation computes a new sparse matrix P whose columns,

$$P(:, j) = (I - \omega D^{-1} A) T(:, j), \tag{1.8}$$

are the result of applying one-iteration of relaxation to each column of the tentative prolongator. In practice, we compute all columns of P in a single operation using a specialized sparse matrix-matrix multiplication algorithm.

1.2.6. Galerkin Product. The construction of the sparse Galerkin product $A_{k+1} = R_k A_k P_k$ in Line 6 of Algorithm 1 is typically implemented with two separate sparse matrix-matrix multiplies: i.e., $(RA)P$ or $R(AP)$. Either way, the first product is of the form $[n \times n] * [n \times n_c]$ (or the transpose), while the second product is of the form $[n_c \times n] * [n \times n_c]$.

Efficient sequential sparse matrix-matrix multiplication algorithms are described in [21, 3]. In these methods the Compressed Sparse Row (CSR) format is used, which provides $\mathcal{O}(1)$ indexing of the matrix rows. As a result, the restriction matrix $R_k = P_k^T$ is formed explicitly in CSR format before the Galerkin product is computed.

While the sequential algorithms for sparse matrix-matrix multiplication are efficient, they rely on a large amount of (per thread) temporary storage, and are therefore not suitable for fine-grained parallelism. Specifically, to compute the sparse product $C = A * B$, the sequential methods use $O(N)$ additional storage, where N is the number of columns in C. In contrast, our approach to sparse matrix-matrix multiplication, detailed in Section 3.3, is formulated in terms of highly-scalable parallel primitives with no such limitations. Indeed, our formulation exploits parallelism at the level of individual matrix entries.

1.3. Spectral radius. For the smoothers such as weighted Jacobi or Chebyshev, a calculation of the spectral radius of a matrix, i.e., the eigenvalue of the matrix with maximum modulus, is often needed in order to yield effective smoothing properties. These smoothers are central to the relaxation scheme in the cycling phase and the prolongation in the setup phase, so we consider the computational impact of these calculations. In our case, an approximation of the spectral radius of $D^{-1}A$ where D is matrix of the diagonal of A.

For any non-Hermitian matrix A, Arnoldi iteration is a method used to reduce A to an upper Hessenberg matrix H by similarity transformations. The eigenvalues of the small fixed-size dense matrix H are then computed directly. This is the approach that we consider in this paper, since it produces accurate estimates of the spectral radius. In Section 5, we see that computing the spectral radius has a non-trivial role in the total cost of the setup phase.

1.4. Cycling Phase. The multigrid cycling or solve phase is detailed in Algorithm 3. Several computations are executed at each level in the algorithm, but as we see in Lines 2-6, the operations are largely sparse matrix-vector multiplications (SpMV). Consequently, on a per-level basis, we observe a strong relationship between the performance of the SpMV and the computations in Lines 2-6. For example, the smoothing sweeps on Lines 2 and 6 are both implemented as affine operations such as $x_k \leftarrow x_k - \omega D^{-1}(Ax_k - b)$, in the case of weighted Jacobi. This is a highly parallelized `AXPY` operation as well as a SpMV. This is also the case for the residual on Line 3, and the restriction and interpolation operations on Lines 4 and 5. Finally, we coarsen to only a few points so that the coarse solve on Line 1 is efficiently handled by either relaxation or a call to CUBLAS [14].

While the computations in the solve phase are straightforward at a high-level, they rely heavily on the ability to perform *unstructured* SpMV operations efficiently. Indeed, even though fine level matrix may exhibit a well-defined structure, e.g., a banded matrix, the unstructured nature of the aggregation process produces coarse-level matrices with less structure. In Section 5 we examine the cost of the solve phase in more detail.

Algorithm 3: AMG Solve: `solve`

parameters: A_k, sparse problem matrix
R_k, restriction matrix
P_k, interpolation matrix
x_k, solution vector
b_k, right-hand-side vector
return: x_k, solution vector

if $k = M$
1 solve $A_k x_k = b_k$
else
2 $x_k \leftarrow \texttt{presmooth}(A_k, x_k, b_k, \mu_1)$ {smooth μ_1 times on $A_k x_k = b_k$}
3 $r_k \leftarrow b_k - A_k x_k$ {compute residual}
4 $r_{k+1} \leftarrow R_k r_r$ {restrict residual}
$e_{k+1} \leftarrow \texttt{solve}(A_{k+1}, R_{k+1}, P_{k+1}, e_{k+1}, r_{k+1})$ {coarse-grid solve}
5 $e_k \leftarrow P_k e_{k+1}$ {interpolate solution}
$x_k \leftarrow x_k + e_k$ {correct solution}
6 $x_k \leftarrow \texttt{postsmooth}(A_k, x_k, b_k, \mu_2)$ {smooth μ_2 times on $A_k x_k = b_k$}

2. Parallel Primitives. Our method for exposing fine-grained parallelism in AMG leverages (data) *parallel primitives* [8, 33]. We use the term primitives to refer to a collection of fundamental algorithms that emerge in numerous contexts such as reduction, parallel prefix-sum (or scan), and sorting. In short, these primitives are to general-purpose computations what BLAS [27] is to computations in linear algebra. Like BLAS, the practice of distilling complex algorithms into parallel primitives has several advantages versus authoring ad hoc codes:
productivity: programming with primitives requires less developer effort,
performance: primitives can be optimized for the target architecture independently,
portability: applications can be ported to a new architecture more easily.
Given the broad scope of their usage, special emphasis has been placed on the performance of primitives and very highly-optimized implementations are readily available for the GPU [33, 29, 30]. The efficiency of our solver, and hence the underlying parallel primitives, is demonstrated in Section 5.

Our AMG solver is implemented almost exclusively with the parallel primitives provided by the Thrust library [24]. Following the model of the C++ Standard Template Library (STL), the algorithms provided by Thrust are *generic* and accept general data types and operators. In the remaining part of this section we identify a few of the most important Thrust algorithms and illustrate their usage. For ease of exposition we omit some of the precise usage details, however the essential characteristics are preserved.

2.1. Reduction. A critical component in many of the computations that we encounter is that of simplifying an array to a single value, or a *reduction*. Reduction is implemented in Thrust with the `reduce` algorithm. By default, the `reduce` algorithm reduces a range of numbers to a single value by successively summing values together:

$$\texttt{reduce}([3, 4, 1, 5, 2]) \rightarrow 15.$$

The same algorithm can be used to determine the maximum entry, by specifying `maximum` for the reduction operator:

$$\texttt{reduce}([3,4,1,5,2],\texttt{maximum}) \to 5.$$

In general, any function that is both commutative and associative is a valid reduction operator.

Reductions involving two vectors are implemented with the `inner_product` algorithm. As the name suggests, by default the algorithm computes the mathematical inner-product (or dot product),

$$\texttt{inner\_product}([3,4,1,5],[0,1,2,3]) \to 21.$$

As with `reduce`, the `inner_product` algorithm supports different reduction and elementwise vector operations, in place of addition and multiplication respectively. For example,

$$\texttt{inner\_product}([1,2,3],[1,2,3],\texttt{logical\_and},\texttt{equals}) \to \texttt{true},$$
$$\texttt{inner\_product}([1,2,3],[1,0,3],\texttt{logical\_and},\texttt{equals}) \to \texttt{false},$$

tests two vectors for equality.

2.2. Scan. Similarly, the parallel prefix-sum, or *scan*, primitive computes the cumulative sum of an array and is a fundamental component of common algorithms such as stream compaction. In Thrust, the `inclusive_scan` algorithm computes the "inclusive" variant of the scan primitive,

$$\texttt{inclusive\_scan}([3,4,1,5,2]) \to [3,7,8,13,15],$$

while the `exclusive_scan` algorithm computes the "exclusive" variant,

$$\texttt{exclusive\_scan}([3,4,1,5,2],10) \to [10,13,17,18,23],$$

which incorporates a user-specified starting value and excludes the final sum. As with reduction, the scan algorithms accept other binary operations such as `maximum`,

$$\texttt{inclusive\_scan}([3,4,1,5,2],\texttt{maximum}) \to [3,4,4,5,5],$$

provided that the operator is associative.

2.3. Transformations. We also encounter *transformations* on arrays, or elementwise map operations, which are implemented with the `transform` algorithm. Unary transformations apply a unary functor to each element of an input array and write the output to another array. For example, transforming an array with the `negate` functor,

$$\texttt{transform}([3,4,1],\texttt{negate}) \to [-3,-4,-1],$$

implements vector negation. In the same way, binary transformations apply a binary functor to pairs of elements in two input arrays and write the output to another array. For example, transforming a pair of vectors with the `plus` functor,

$$\texttt{transform}([3,4,1],[4,5,7],\texttt{plus}) \to [7,9,8],$$

implements vector addition.

2.4. Gathering and Scattering. Related to transformation are the `gather` and `scatter` algorithms,

$$\texttt{gather}([3,0,2],[11,12,13,14]) \rightarrow [14,11,13],$$
$$\texttt{scatter}([3,0,2],[11,12,13],[*,*,*,*]) \rightarrow [12,*,13,11],$$

which copy values based on an index map ($[3,0,2]$ in the examples). Here, the placeholder $*$ represents elements of the output range that are not changed by the algorithm. Whenever an algorithm requires an indirect load (e.g., $v = X[map[i]]$) or store (e.g., $Y[map[i]] = v$) operation, a `gather` and `scatter` is typically required. Predicated versions of the functions,

$$\texttt{gather\_if}([3,0,2],[true,false,true],[11,12,13,14],[*,*,*]) \rightarrow [14,*,13],$$
$$\texttt{scatter\_if}([3,0,2],[true,false,true],[11,12,13],[*,*,*,*]) \rightarrow [*,*,13,11],$$

copy values conditionally.

2.5. Stream Compaction. It is often necessary to filter elements of an array based on a given condition or *predicate*. This process, known as stream compaction [33], is implemented with Thrust by combining the `copy_if` algorithm with a predicate functor, and is invaluable in expressing certain components of AMG such as strength-of-connection. For example,

$$\texttt{copy\_if}([3,4,1,5,2],\texttt{is\_even}) \rightarrow [4,2],$$

copies the even elements of an array to another array, preserving their original order. Here `is_even` is a unary functor that returns `true` if the argument is even and `false` otherwise.

The first variant of `copy_if`, applies the predicate functor to the values of the input sequence and copies a subset of them to the output range. A second variant applies the predicate to a separate sequence and copies a subset of the corresponding values to the output. For example,

$$\texttt{copy\_if}([0,1,2,3,4],[3,4,1,5,2],\texttt{is\_even}) \rightarrow [1,4],$$

outputs the positions of even elements, as opposed to the values of even elements.

2.6. Segmented Reduction. The `reduce` algorithm reduces an entire array, which is typically large, down to a single value. In practice it is often necessary to compute a reduction for many different arrays at once, however using `reduce` for each one is impractical since the array size is generally too small to expose sufficient fine-grained parallelism. For this reason Thrust provides a `reduce_by_key` algorithm that implements *segmented* reduction. In segmented reduction, an array of keys determines which values belong to a segment. For example, the key array $[0,0,1,1,1,2]$ defines three segments of lengths two, three, and one respectively. When the `reduce_by_key` algorithm encounters a set of adjacent keys that are equivalent, it reduces the corresponding values together and writes the key and the reduced value to separate output arrays. For example,

$$\texttt{reduce\_by\_key}([0,0,1,1,1,2],[10,20,30,40,50,60]) \rightarrow [0,1,2],[30,120,60].$$

Note that the key and value sequences are stored in separate arrays. This "structure of arrays" representation is generally more computationally efficient than the alternative "array of structures" representation, where keys and values are interleaved in memory.

2.7. Sorting. Sorting is an effective primitive whenever disordered data must be binned or easily indexed. This is helpful in many of our transformations in the AMG setup phase. By default, the `sort` algorithm sorts data in ascending order,

$$\texttt{sort}([3,4,1,5,2]) \rightarrow [1,2,3,4,5],$$

which is equivalent to specifying that elements should be compared using the standard `less` comparison functor

$$\texttt{sort}([3,4,1,5,2], \texttt{less}) \rightarrow [1,2,3,4,5].$$

Replacing the `less` comparison operator with `greater` sorts data in descending order,

$$\texttt{sort}([3,4,1,5,2], \texttt{greater}) \rightarrow [5,4,3,2,1].$$

In general, any valid strict weak ordering can be used to compare values.

Thrust also provides the `sort_by_key` algorithm for sorting (logical) key-value pairs. In the key-value sort, the keys are sorted as usual, while the value corresponding to each key is moved to the same position in the values array,

$$\texttt{sort\_by\_key}([3,4,1,5,2],[10,20,30,40,50]) \rightarrow [1,2,3,4,5],[30,50,10,20,40].$$

Combining the `sort_by_key` algorithm with the `reduce_by_key` algorithm facilitates MapReduce-style [15] computations. For example, sorting the key-value pairs by key,

$$\texttt{sort\_by\_key}([0,1,0,1,0],[10,20,30,40,50]) \rightarrow [0,0,0,1,1],[10,30,50,20,40],$$

places all values with the same key into contiguous segments, allowing them to be summed together,

$$\texttt{reduce\_by\_key}([0,0,0,1,1],[10,30,50,20,40]) \rightarrow [0,1],[90,60],$$

using segmented reduction.

The `sort` and `sort_by_key` algorithms make no guarantees about the stability of the sorting process. Specifically, if two (or more) keys are equivalent, meaning neither is less than the other, then these algorithms are free to reorder such keys in the process of sorting the data. In contrast, a *stable* sorting algorithm, such as `stable_sort` and `stable_sort_by_key`, preserves the relative ordering of equivalent keys. In practice, stable sorting methods are useful when reordering a set of values based on multiple key vectors, for instance reordering the values based on both the row and column indices of a matrix.

3. Parallel Hierarchy Construction. In this section we describe a method for adapting the multigrid setup phase of Algorithm 1. We focus on expressing the components of the setup phase so that they exploit the efficiency of several algorithms on the GPU. The main feature of such algorithms is that they expose very fine-grained parallelism. Complete source code for the methods presented here are available in the open-source Cusp library [6]. Our primary contributions are the parallel aggregation and sparse matrix-matrix multiplication algorithms. Although we have restricted our attention to a particular AMG method, namely smoothed aggregation [37], the algorithms discussed here are broadly applicable as other AMG variants share many of the same basic components.

The methods described in this section are designed for the *coordinate* (COO) sparse matrix format. The COO format is comprised of three arrays `I`, `J`, and `V`, which store the row indices, column indices, and values, respectively, of the matrix entries. We further assume that the COO entries are sorted by ascending row index. For example, the matrix

$$\mathtt{A} = \begin{bmatrix} 2 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 2 \end{bmatrix},$$

is represented in coordinate format by three arrays

$$\begin{aligned} \mathtt{I} &= \begin{bmatrix} 0 & 0 & 1 & 1 & 1 & 2 & 2 & 2 & 3 & 3 \end{bmatrix}, \\ \mathtt{J} &= \begin{bmatrix} 0 & 1 & 0 & 1 & 2 & 1 & 2 & 3 & 2 & 3 \end{bmatrix}, \\ \mathtt{V} &= \begin{bmatrix} 2 & -1 & -1 & 2 & -1 & -1 & 2 & -1 & -1 & 2 \end{bmatrix}, \end{aligned}$$

of ten entries each. Although the COO representation is not generally the most efficient sparse matrix storage format, it is simple and easily manipulated.

3.1. Strength of Connection. The symmetric strength-of-connection algorithm (cf. Section 1.2.2) is straightforward to implement using parallel primitives. Given a matrix A in coordinate format we extract the matrix diagonal by comparing the row index array (`I`) to the column index array (`J`),

$$\begin{aligned} \mathtt{D} &= [0,0,0,0], \\ \mathtt{is\_diagonal} &= \mathtt{transform}(\mathtt{I},\mathtt{J},\mathtt{equals}), \\ &= [\mathtt{true},\mathtt{false},\mathtt{false},\mathtt{true},\mathtt{false},\mathtt{false},\mathtt{true},\mathtt{false},\mathtt{false},\mathtt{true}], \\ \mathtt{D} &= \mathtt{scatter\_if}(\mathtt{V},\mathtt{I},\mathtt{is\_diagonal},\mathtt{D}), \\ &= [2,2,2,2], \end{aligned}$$

and scattering the corresponding entries in the values array (`V`) when they are equal.

The second loop in Algorithm 2 is implemented with stream compaction. First we obtain arrays containing `D[i]` and `D[j]` for each index in `I` and `J` respectively,

$$\begin{aligned} \mathtt{Di} &= \mathtt{gather}(\mathtt{I},\mathtt{D}), \\ &= [2,2,2,2,2,2,2,2,2,2], \\ \mathtt{Dj} &= \mathtt{gather}(\mathtt{J},\mathtt{D}), \\ &= [2,2,2,2,2,2,2,2,2,2], \end{aligned}$$

from which the strength-of-connection threshold (i.e., $\theta\sqrt{|D(I_n)| \cdot |D(J_n)|}$ in Algorithm 2) of each entry is computed:

$$\mathtt{threshold} = \mathtt{transform}(\mathtt{Di},\mathtt{Dj},\mathtt{soc\_threshold(theta)}),$$

using a specially-defined functor `soc_threshold`. Next, each entry in the values array (`V`) is tested against the corresponding threshold to determine which entries are strong,

$$\begin{aligned} \mathtt{is\_strong} &= \mathtt{transform}(\mathtt{V},\mathtt{threshold},\mathtt{greater}), \\ &= [\mathtt{true},\mathtt{true},\mathtt{true},\mathtt{true},\mathtt{true},\mathtt{true},\mathtt{true},\mathtt{true},\mathtt{true},\mathtt{true}]. \end{aligned}$$

Finally, the matrix entries corresponding to strong connections are determined using stream compaction to form a coordinate representation of the strength-of-connection matrix, namely

$$\begin{aligned}\texttt{Ci} &= \texttt{copy\_if}(\texttt{I}, \texttt{is\_strong}, \texttt{identity}),\\ \texttt{Cj} &= \texttt{copy\_if}(\texttt{J}, \texttt{is\_strong}, \texttt{identity}),\\ \texttt{Cv} &= \texttt{copy\_if}(\texttt{V}, \texttt{is\_strong}, \texttt{identity}).\end{aligned}$$

In practice we do not construct arrays such as `Di` and `threshold` explicitly in memory. Instead, the `gather` and `transform` operations are *fused* with subsequent algorithms to conserve memory bandwidth using Thrust's `permutation_iterator` and `transform_iterator`. Similarly, the three calls to `copy_if` are combined into a single stream compaction operation using a `zip_iterator`. The translation from the explicit version of the algorithm (shown here) and the more efficient, fused version (not shown) is only a mechanical transformation. We refer to the Cusp [6] source code for additional detail.

3.2. Aggregation. The sequential aggregation method, which is described in Section 1.2.3 is designed to produce aggregates with a particular structure. Unfortunately, the greedy nature of the algorithm introduces sequential dependencies that prevent a direct parallelization of the algorithm's first phase. In this section we describe a parallel analog of the sequential method that produces aggregates with the same essential properties. Our method is based on a generalized *maximal independent set* algorithm and is amenable to fine-grained parallelism. A similar parallel aggregation strategy is described in [35], albeit with a different implementation.

There are two observations regarding the aggregation depicted in Figure 1.1 that lead to the description of our aggregation method. First, no two root nodes of the aggregates are within a distance of two edges apart. Second, if any unaggregated node is separated from all existing root nodes by more than two edges then it is free to become the root of a new aggregate. Taken together, these conditions define a collection of root nodes that are a distance-2 maximal independent set, which we denote MIS(2) and formalize in Definition 3.1. The first property ensures independence of the nodes — i.e., that the minimum distance between any pair of root nodes exceeds a given threshold. The second property ensures maximality — i.e., no additional node can be added to the set without violating the property of independence. The standard definition of a maximal independent set, which we denote MIS(1), is consistent with this definition except with a distance of one. We defer a complete description of the generalized maximal independent set algorithm to the Appendix. For the remainder of this section, we assume that a valid MIS(2) is computed efficiently in parallel.

DEFINITION 3.1 (MIS(k)). *Given a graph $G = (V, E)$, let $V_{root} \subset V$ be a set of root nodes, and denote by $d_G(\cdot,\cdot)$ the distance or shortest path between two vertices in the graph. Then, V_{root} is a maximal independent set of distance k, or MIS(k) if the following hold:*

1. *(independence) Given any two vertices $u, v \in V_{root}$, then $d_G(u, v) > k$.*
2. *(maximality) There does not exist $u \in V \setminus V_{root}$ such that $d_G(u, v) > k$ for all $v \in V_{root}$.*

Given a valid MIS(2) the construction of aggregates is straightforward. Assume that the MIS(2) is specified by an array of values in $\{0, 1\}$, where a value of 1 indicates that corresponding node is a member of the independent set and value of 0 indicates otherwise. We first enumerate the MIS(2) nodes with an `exclusive_scan` operation,

giving each set node a unique index in $[0, N)$, where N is the number of nodes in the set. For example, on a linear graph with ten nodes, a MIS(2) with four set nodes,

$$\texttt{mis} = [1, 0, 0, 1, 0, 0, 1, 0, 0, 1],$$

is enumerated as

$$\begin{aligned}\texttt{enum} &= \texttt{exclusive\_scan}(\texttt{mis}, 0),\\ &= [0, 1, 1, 1, 2, 2, 2, 3, 3, 3].\end{aligned}$$

Since the i-th node in the independent set serves as the root of the i-th aggregate, the only remaining task it so propagate the root indices outward.

The root indices are communicated to their neighbors with two operations resembling a sparse matrix-vector multiply, $y = Ax$. Conceptually, each unaggregated node looks at neighbors for the index of an aggregate. In the first step, all neighbors of a root node receive the root node's aggregate index — i.e., the value resulting from the `exclusive_scan` operation. In the second step, the remaining unaggregated nodes receive the aggregate indices of their neighbors, at least one of which must belong to an aggregate. As before, in the presence of multiple neighboring aggregates a selection is made arbitrarily. The two sparse matrix-vector operations are analogous to the first and second phases of the sequential algorithm respectively (cf. Section 1.2.3). Our implementation of the parallel aggregate propagation step closely follows the existing techniques for sparse matrix-vector multiplication [5].

Figure 3.1 depicts a MIS(2) for a regular mesh and the corresponding aggregates rooted at each independent set node. Although the root nodes are selected by a randomized process (see the Appendix) the resulting aggregates are qualitatively similar to those chosen by the sequential algorithm in Figure 1.1a. Indeed, with the appropriate permutation of graph nodes, the sequential aggregation algorithm selects the same root nodes as the randomized MIS(2) method.

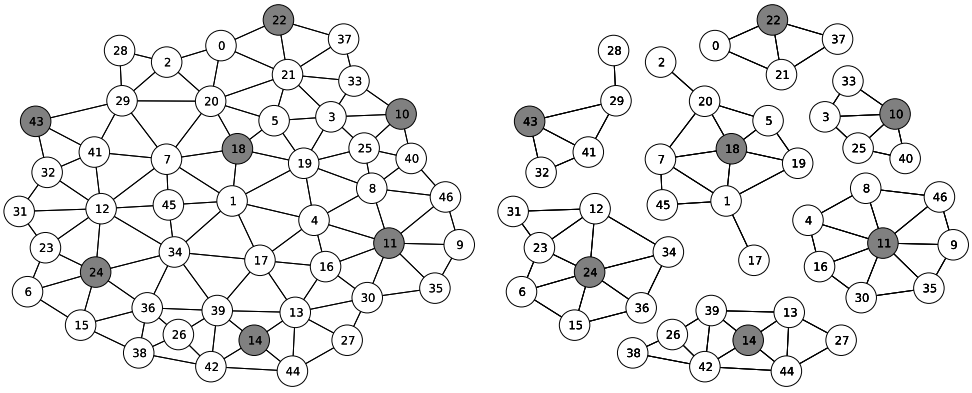

(a) MIS(2) (b) Aggregates

FIG. 3.1. *Parallel aggregation begins with a MIS(2) set of nodes, colored in gray, and represent the root of an aggregate — e.g., node 18. As in the sequential method, nodes adjacent to a root node are incorporated into the root node's aggregate in the first phase. In the second phase, unaggregated nodes join an adjacent aggregate — e.g., nodes 12, 19, and 24 for root node 18.*

3.3. Sparse Matrix-Matrix Multiplication. Efficiently computing the product $C = AB$ of sparse matrices A and B is challenging, especially in parallel. The central difficulty is that, for irregular matrices, the structure of the output matrix C has a complex and unpredictable dependency on the input matrices A and B. Indeed, simply computing the sparsity pattern of the output is not an efficient or straightforward algorithm.

The sequential sparse matrix-matrix multiplication algorithms mentioned in Section 1.2.6 do not admit immediate, fine-grained parallelization. Specifically, for matrices A and B of size $[k \times m]$ and $[m \times n]$ the methods requires $O(n)$ temporary storage to determine the entries of each sparse row in the output. As a result, a straightforward parallelization of the sequential scheme requires $O(n)$ storage per thread, which is untenable when seeking to develop tens of thousands of independent threads of execution. While it is possible to construct variations of the sequential method with lower per-thread storage requirements, any method that operates on the granularity of matrix rows — i.e., distributing matrix rows over threads — requires a non-trivial amount of per-thread state and suffers load imbalances for certain input. As a result, we have designed an algorithm for sparse matrix-matrix multiplication based on sorting, segmented reduction, and other operations which are well-suited to fine-grained parallelism as discussed in Section 2.

As an example, we demonstrate our algorithm for computing $\mathtt{C} = \mathtt{A} * \mathtt{B}$, where

$$\mathtt{A} = \begin{bmatrix} 5 & 10 & 0 \\ 15 & 0 & 20 \end{bmatrix} \quad \text{and} \quad \mathtt{B} = \begin{bmatrix} 25 & 0 & 30 \\ 0 & 35 & 40 \\ 45 & 0 & 50 \end{bmatrix}, \tag{3.1}$$

have 4 and 6 nonzero entries respectively. The matrices are stored in coordinate format as

$$\mathtt{A} = \begin{bmatrix} (0,0,\ 5) \\ (0,1,10) \\ (1,0,15) \\ (1,2,20) \end{bmatrix} \quad \text{and} \quad \mathtt{B} = \begin{bmatrix} (0,0,25) \\ (0,2,30) \\ (1,1,35) \\ (1,2,40) \\ (2,0,45) \\ (2,2,50) \end{bmatrix}, \tag{3.2}$$

where each (i, j, v) tuple of elements represents the row index, column index, and value of the matrix entry. We note that the (i, j, v) tuples are only a logical construction used to explain the algorithm. In practice the coordinate format is stored in a "structure of arrays" layout with three separate arrays.

To expose fine-grained parallelism, our parallel sparse matrix-matrix multiplication algorithm proceeds in three stages:

1. *Expansion* of $\mathtt{A} * \mathtt{B}$ into a intermediate coordinate format $\mathtt{T}$
2. *Sorting* of $\mathtt{T}$ by row and column indices to form $\hat{\mathtt{T}}$
3. *Compression* of $\hat{\mathtt{T}}$ by summing duplicate values for each matrix entry

In the first stage, $\mathtt{T}$ is formed by multiplying each entry $\mathtt{A}(i, j)$ of the left matrix with

every entry in row j of $\mathtt{B}$, $\mathtt{B}(j,k)$ for all k. In our example, the intermediate format

$$\mathtt{T} = \begin{bmatrix} (0,0,\ 125) \\ (0,2,\ 150) \\ (0,1,\ 350) \\ (0,2,\ 400) \\ (1,0,\ 375) \\ (1,2,\ 450) \\ (1,0,\ 900) \\ (1,2,1000) \end{bmatrix}, \tag{3.3}$$

contains a total of 8 entries, including duplicate entries for the coordinates $(0,2)$, $(1,0)$, and $(1,2)$. The first two entries of $\mathtt{T}$ result from multiplying the first entry of $\mathtt{A}$, $\mathtt{A}(0,0)$, with the entries in the corresponding row of $\mathtt{B}$, namely $\mathtt{B}(0,0)$ and $\mathtt{B}(0,2)$. The next two entries of $\mathtt{T}$ result from multiplying the second entry of $\mathtt{A}$, $\mathtt{A}(0,1)$, with the entries in the corresponding row of $\mathtt{B}$, namely $\mathtt{B}(1,1)$ and $\mathtt{B}(1,2)$, and so on. The expansion stage is implemented in parallel using gather, scatter, and prefix-sum operations. We refer to the Cusp source code [6] for further detail.

The result of the expansion phase is an intermediate coordinate format with possible duplicate entries that is sorted by row index, but not by column index. The second stage of the sparse matrix-matrix multiplication algorithm sorts the entries of $\mathtt{T}$ into lexicographical order. For example, sorting the entries of $\mathtt{T}$ by the (i,j) coordinate yields

$$\hat{\mathtt{T}} = \begin{bmatrix} (0,0,\ 125) \\ (0,1,\ 350) \\ (0,2,\ 150) \\ (0,2,\ 400) \\ (1,0,\ 375) \\ (1,0,\ 900) \\ (1,2,\ 450) \\ (1,2,1000) \end{bmatrix}, \tag{3.4}$$

from which the duplicate entries at coordinates $(0,2)$, $(1,0)$, and $(1,2)$ are easily identified. The lexicographical reordering is efficiently implemented with Thrust's `stable_sort_by_key` algorithm.

The final stage of the algorithm compresses duplicate coordinate entries while summing their corresponding values. Since $\hat{\mathtt{T}}$ is sorted by row and column, the duplicate entries are stored contiguously and are compressed with a single application of the `reduce_by_key` algorithm. In our example $\hat{\mathtt{T}}$ becomes

$$\mathtt{C} = \begin{bmatrix} (0,0,\ 125) \\ (0,1,\ 350) \\ (0,2,\ 550) \\ (1,0,1275) \\ (1,2,1450) \end{bmatrix}, \tag{3.5}$$

where each of the duplicate entries $\mathtt{C}(0,2) = 150 + 400$, $\mathtt{C}(1,0) = 375 + 900$, and $\mathtt{C}(1,2) = 450 + 1000$ have been combined. The compressed result is now a valid, duplicate-free coordinate representation of the matrix

$$\mathtt{A} * \mathtt{B} = \begin{bmatrix} 125 & 350 & 550 \\ 1275 & 0 & 1450 \end{bmatrix}.$$

All three stages of our sparse matrix-matrix multiplication algorithm expose ample fine-grained parallelism. Indeed, even modest problem sizes generate enough independent threads of execution to fully saturate the GPU. Furthermore, because we have completely flattened the computation into efficient data-parallel algorithms — e.g., `gather`, `scatter`, `scan`, `stable_sort_by_key`, etc. — our implementation is insensitive to the (inherent) irregularity of sparse matrix-matrix multiplication. As a result, even pathological inputs do not create substantial imbalances in the work distribution among threads.

Another benefit is that our parallel sparse matrix-matrix algorithm has the same computational complexity as the sequential method, which is $O(nnz(\mathtt{T}))$, the number of entries in our intermediate format. Therefore our method is "work-efficient" [8], since (1) the complexity of the sequential method is proportional to the size of the intermediate format `T` and (2) the work involved at each stage of the algorithm is linear in the size of `T`. The latter claim is valid since Thrust employs work-efficient implementations of parallel primitives such as `scan`, `reduce_by_key` and `stable_sort_by_key`.

One practical limitation of the method as described above is that the memory required to store the intermediate format is potentially large. For instance, if `A` and `B` are both square, $n \times n$ matrices with exactly K entries per row, then $O(nK^2)$ bytes of memory are needed to store `T`. Since the input matrices are generally large themselves ($O(nK)$ bytes) it is not always possible to store a K-times larger intermediate result in memory. In the limit, if `A` and `B` are dense matrices (stored in sparse format) then $O(n^3)$ storage is required. As a result, our implementation allocates a workspace of bounded size and decomposes the matrix-matrix multiplication $\mathtt{C} = \mathtt{A} * \mathtt{B}$ into several smaller operations $\mathtt{C(slice,:)} = \mathtt{A(slice,:)} * \mathtt{B}$, where *slice* is a subset of the rows of `A`. The final result is obtained by simply concatenating the coordinate representations of all the partial results together $\mathtt{C} = [\mathtt{C(slice\_0,:)}, \mathtt{C(slice\_1,:)}, \ldots]$. In practice this sub-slicing technique introduces little overhead because the workspace is still large enough to fully utilize the device. We refer to the Cusp source code [6] for additional details.

3.4. Prolongation and Restriction. The tentative prolongation, prolongation smoothing, and restriction construction steps of Algorithm 1 (Lines 3, 4, and 5), are relatively straightforward to express with parallel primitives. The tentative prolongation operation is constructed by gathering the appropriate entries from the near-nullspace vectors stored in B according to the sparsity pattern defined by Agg. Then, the columns are normalized, first by transposing the matrix, which has the effect of sorting the matrix entries by column index, and then computing the norm of each column using the `reduce_by_key` algorithm. Specifically, the transpose of a coordinate format matrix such as

$$\begin{aligned}
\mathtt{I} &= \begin{bmatrix} 0 & 1 & 2 & 3 & 4 & 5 \end{bmatrix}, \\
\mathtt{J} &= \begin{bmatrix} 0 & 1 & 1 & 0 & 1 & 0 \end{bmatrix}, \\
\mathtt{V} &= \begin{bmatrix} 0 & 0 & 1 & 1 & 2 & 2 \end{bmatrix},
\end{aligned}$$

is computed with by reordering the column indices of the matrix, and applying the same permutation to the rows and values

$$\begin{aligned}
\texttt{TransI}, \texttt{Permutation} &= \texttt{stable_sort_by_key}(\texttt{J}, [0,1,2,3,4]) \\
&= \begin{bmatrix}0 & 0 & 0 & 1 & 1 & 1\end{bmatrix}, \begin{bmatrix}0 & 3 & 5 & 1 & 2 & 4\end{bmatrix}, \\
\texttt{TransJ} &= \texttt{gather}(\texttt{Permutation}, \texttt{I}) \\
&= \begin{bmatrix}0 & 3 & 5 & 1 & 2 & 4\end{bmatrix}, \\
\texttt{TransV} &= \texttt{gather}(\texttt{Permutation}, \texttt{V}) \\
&= \begin{bmatrix}0 & 1 & 2 & 0 & 1 & 2\end{bmatrix}.
\end{aligned}$$

Then the squares of the transposed values array are calculated, followed by row sums,

$$\begin{aligned}
\texttt{Squares} &= \texttt{tranform}(\texttt{TransV}, \texttt{TransV}, \texttt{multiplies}) \\
&= \begin{bmatrix}0 & 1 & 4 & 0 & 1 & 4\end{bmatrix}, \\
\texttt{Rows}, \texttt{Sums} &= \texttt{reduce_by_key}(\texttt{TransI}, \texttt{Squares}, \texttt{multiplies}) \\
&= \begin{bmatrix}0 & 1\end{bmatrix}, \begin{bmatrix}9 & 9\end{bmatrix},
\end{aligned}$$

which correspond to column sums of the original matrix.

Next, the final prolongator P is obtained by smoothing the columns of T according to the formula in Section 1.2.5. Here, we apply a specialized form of the general sparse matrix-matrix multiplication scheme described in Section 3.3. Specifically, we exploit the special structure of the tentative prolongator, whose rows contain at most one nonzero entry, when computing the expression AT.

Finally, the transpose of the prolongation operator is calculated explicitly $R = P^T$, and Galerkin triple-matrix product $A_{k+1} = R_k(A_k P_k)$ is computed with two separate sparse matrix-matrix multiplies. As mentioned above, the transpose operation is fast, particularly for the COO format. In Section 5 we show that the transpose is very efficient on the GPU.

4. Parallel Multigrid Cycling. After the multigrid hierarchy has been constructed using the techniques in Section 3, the cycling of Algorithm 3 proceeds. In this section we describe the components of the multigrid cycling and how they are parallelized on the GPU.

4.1. Vector Operations. The multigrid cycle involves several standard vector operations. In Algorithm 3, the residual vector computation and the coarse grid correction steps require vector-vector subtraction and addition respectively. While these operations could be fused with the preceding sparse matrix-vector products for potential efficiency, or could be carried out with `DAXPY` in CUBLAS [14], we have implemented equivalent routines with Thrust's `transform` algorithm to keep Cusp self-contained. Similarly, the vector norms (`DNRM2`) and inner products (`DDOT`) that arise in multigrid cycling have been implemented with `reduce` and `inner_product` in Thrust respectively.

4.2. Sparse Matrix-Vector Multiplication. Sparse matrix-vector multiplication (SpMV), which involves (potentially) irregular data structures and memory access patterns, is more challenging to implement than the aforementioned vector operations. Nevertheless efficient techniques exist for matrices with a wide variety of sparsity patterns [11, 9, 38, 16, 39, 4, 5]. Our implementations of sparse matrix-vector multiplication are described in [5, 7]. In Algorithm 3 sparse matrix-vector

multiplication is used to compute the residual, to restrict the residual, to interpolate the coarse-level solution onto the finer grid, and in many cases, to implement the pre- and post-smoother.

In Section 3 we describe a method for constructing the AMG hierarchy in parallel using the coordinate (COO) sparse matrix format. The COO format is simple to construct and manipulate, and therefore is well-suited for the computations that arise in the setup phase. However, the COO format is generally not the most efficient for the SpMV operations [7]. Fortunately, once the hierarchy is constructed it remains unchanged during the cycling phase. As a result, it is beneficial to convert the sparse matrices stored throughout the hierarchy to an alternative format that achieves higher SpMV performance.

In Cusp, conversions between COO and other sparse matrix formats, such as CSR, DIA, ELL, and HYB [7], are inexpensive, as shown in Table 4.1. Here we see that the conversion from COO to CSR is trivial, while to more structured formats such as ELL and HYB is of minimal expense (*note: the conversion to itself represents a straight copy*). When reporting performance figures in Section 5 we include the COO to HYB conversion time in the setup phase.

From\To	COO	CSR	ELL	HYB
COO	5.09	6.35	20.10	23.55
CSR	7.80	4.03	21.61	24.84
ELL	17.02	18.26	5.90	22.32
HYB	63.27	69.40	83.08	4.12

TABLE 4.1
Sparse matrix conversion times (ms) for an unstructured mesh with $1M$ vertices and $8M$ edges.

4.3. Smoothing. Gauss-Seidel relaxation is a popular multigrid smoother with several attractive properties. For instance, the method requires only $O(1)$ temporary storage and converges for any symmetric, positive-definite matrix. Unfortunately, the standard Gauss-Seidel scheme does not admit an efficient parallel implementation. Jacobi relaxation is a simple and highly parallel alternative to Gauss-Seidel, albeit without the same computational properties. Whereas Gauss-Seidel updates each unknown immediately, Jacobi's method updates all unknowns in parallel, and therefore requires a temporary vector. Additionally, a weighting or damping parameter ω must be incorporated into Jacobi's method to ensure convergence. The weighted Jacobi method is written in matrix form as,

$$I - \frac{\omega}{\rho(D^{-1}A)} D^{-1}A$$

where D is a matrix containing the diagonal elements of A. Since the expression is simply a sparse matrix-vector multiply and an efficient vector operation, Jacobi's method exposes substantial fine-grained parallelism.

We note that sparse matrix-vector multiplication is the main workhorse for several other highly-parallel smoothers. Indeed, so-called polynomial smoothers

$$x \leftarrow x + P(A)r, \tag{4.1}$$

where $P(A)$ is a polynomial in matrix A, are almost entirely implemented with sparse matrix-vector products. In this case, the coefficients of the polynomial are often chosen

to optimize some property of the smoother. For instance, Chebyshev polynomials are designed to have minimal amplitude over a given interval, allowing for optimal damping of specific eigenmodes. We refer to [1] for a comprehensive treatment of parallel smoothers and their associated tradeoffs.

5. Evaluation. In this section we examine the performance of a GPU implementation of the proposed method. We investigate both the setup phase of Algorithm 1 and the solve phase of Algorithm 3, and find tangible speed-ups in each. We first give an overview of the software and hardware used in the numerical experiments, followed by a dissection of the setup and cycling performance.

5.1. Test Platforms. The specifications of our test system are listed in Table 5.1. Our system is configured with CUDA v4.0 [31] and Thrust v1.4 [24]. As discussed in Section 2, Thrust provides many highly optimized GPU parallel algorithms for reduction, sorting, etc. The entire setup phase, and most of the cycling phase, of our GPU method is implemented with Thrust.

Testbed	
GPU	NVIDIA Tesla C2050
CPU	Intel Core i7 950
CLOCK	3.07 GHz
OS	Ubuntu 10.10
Host Compiler	GCC 4.4.5
Device Compiler	NVCC 4.0

TABLE 5.1
Specifications of the test platform.

As a basis for comparison, we also consider the Intel Math Kernel Library (MKL). MKL contains many sparse BLAS subroutines such as matrix-vector and matrix-matrix multiplication. In our performance results we use version 10.3 and link to the multithreaded library.

Trilinos is a software framework provided by Sandia National Laboratories for the solution of large-scale, complex multi-physics engineering simulations on MPI-based clusters. Trilinos includes a smoothed aggregation-based AMG preconditioner for solving large, sparse linear systems in the ML package [18]. In our comparison, we use Trilinos version 10.6 and specifically ML version 5.0 for the solver. The ML results are presented in order to provide context on the performance of our solver in comparison to a well-known software package.

5.2. Model Problem. Table 5.2 describes the sparse matrices considered in our performance evaluation. We present results on both structured and unstructured grids derived from a tessellation of the unit square or unit cube. The unstructured nature of the tessellation is reflected in the sparse matrix, where block and banded patterns are not easily deduced, as they are in the structured case. We consider this case here, since many preconditioners rely on the structured nature of the problem, whereas one advantage of AMG is its relative indifference to matrix structure.

The problem we consider is a 2D and 3D Poisson problem with Dirichlet boundary conditions. Since AMG is known to perform well on such problems, this choice allows us to focus directly on the efficiency of the parallel implementation, rather than on the merits of AMG as a preconditioner.

Matrix	Unknowns	Nonzeros
1a. 2D FD, 5-point	1,048,576	5,238,784
1b. 2D FE, 9-point	1,048,576	9,424,900
2a. 3D FD, 7-point	1,030,301	7,150,901
2b. 3D FE, 27-point	1,030,301	27,270,901
3a. 2D FE, $h \approx 0.03$	550,387	3,847,375
3b. 2D FE, $h \approx 0.02$	1,182,309	8,268,165
3c. 2D FE, $h \approx 0.015$	2,185,401	15,287,137
4. 3D FE, $h \approx 0.15$	1,088,958	17,095,986

TABLE 5.2
Details of the model problem. Here h represents an average mesh diameter for the tessellation.

5.3. Component Performance. Before considering the entire setup and solve phases, we study the performance of their individual components. Table 5.4 reports timings for sparse matrix-vector multiplication, sparse matrix transposition, and sparse matrix-matrix multiplication using MKL and Cusp for the CPU and GPU results, respectively. All computations use double precision (64-bit) floating point arithmetic. The timings in Tables 5.3 and 5.4 are reported in milliseconds, averaged over 100 operations.

Table 5.3 demonstrates the performance of the level 1 BLAS functions `DDOT` and `DAXPY`, which are applied frequently during the cycling phase of the solver. These algorithms execute very few arithmetic operations per memory access and are therefore limited by the available memory bandwidth. For the largest input size the GPU achieves a maximum speedup of 6.62 in `DDOT`. On the smallest input size, the GPU results in a more modest speedup of 2.75 in `DDOT`, which is attributed to fixed costs in our implementation. In contrast, the `DAXPY` speedup is relatively constant across input sizes and simply reflects the relative memory bandwidth of the two processors.

Operation	Size	CPU	GPU	Speedup
`DDOT`	1M	0.88	0.32	2.75
	2M	1.80	0.45	3.98
	4M	3.55	0.71	4.98
	8M	7.13	1.24	5.76
	16M	14.15	2.28	6.20
	32M	28.99	4.38	6.62
`DAXPY`	1M	1.28	0.21	6.08
	2M	2.60	0.42	6.25
	4M	5.22	0.83	6.29
	8M	10.42	1.66	6.29
	16M	20.82	3.33	6.25
	32M	42.54	6.70	6.35

TABLE 5.3
`DDOT` and `DAXPY` times (ms).

Sparse matrix-vector multiplication operations are used heavily in the cycling phase of the solver (cf. Algorithm 3) and to a lesser extent in the setup phase (cf. Algorithm 1) as well. Table 5.4 reports SpMV timings using the Hybrid (HYB) format on the GPU and Compressed Sparse Row (CSR) format on the CPU. As with level 1 BLAS operations, the SpMV operation performs few arithmetic operations per memory access and is memory bound. However, sparse matrix-vector multiplication performance is *also* sensitive to the sparsity structure of the matrix, as it impacts the locality of memory accesses and creates potential variation in the amount of work per thread. Therefore, SpMV does not always saturate memory bandwidth to the same degree as algorithms with simpler memory access patterns such as `DAXPY` (see [5, 7] for a more detailed study of GPU SpMV performance). Across the cases considered, the GPU achieves an average speedup of 6.00 in comparison to the CPU. Figure 5.5 demonstrates that although the CPU SpMV performance does benefit from multi-threading, the marginal speedup diminishes rapidly as the available memory bandwidth is saturated.

Operation	Matrix	CPU	GPU	Speedup
SpMV	1a. 2D FD	6.49	0.78	8.33
	1b. 2D FE	9.42	1.32	7.12
	2a. 3D FD	8.07	0.99	8.12
	2b. 3D FE	22.97	3.44	6.67
	3a. 2D FE	4.18	0.88	4.76
	3b. 2D FE	8.87	1.70	5.20
	3c. 2D FE	16.95	3.19	5.32
	4. 3D FE	16.40	6.46	2.54
Trans.	1a. 2D FD	27.34	9.65	2.83
	1b. 2D FE	26.27	10.05	2.61
	2a. 3D FD	39.73	14.52	2.74
	2b. 3D FE	47.83	16.70	2.86
	3a. 2D FE	15.26	5.79	2.63
	3b. 2D FE	33.93	11.23	3.02
	3c. 2D FE	68.70	20.04	3.43
	4. 3D FE	54.60	17.66	3.09
RAP	1a. 2D FD	247.29	182.89	1.35
	1b. 2D FE	348.84	280.69	1.24
	2a. 3D FD	665.29	425.12	1.56
	2b. 3D FE	1314.16	1271.22	1.03
	3a. 2D FE	254.03	121.82	2.09
	3b. 2D FE	549.75	265.38	2.07
	3c. 2D FE	1032.59	488.75	2.11
	4. 3D FE	1749.95	1025.24	1.71

TABLE 5.4
SpMV time / Transpose time / Galerkin Product time.

A sparse matrix transpose operation is used at each level of in the setup phase (cf. Algorithm 1) to obtain the restriction operator $R_k = P_k^T$ from the prolongation operator P_k. Our parallel algorithm for transpose is 2.90 times faster than MKL on average. We note that the MKL transpose does not benefit from multithreading, sug-

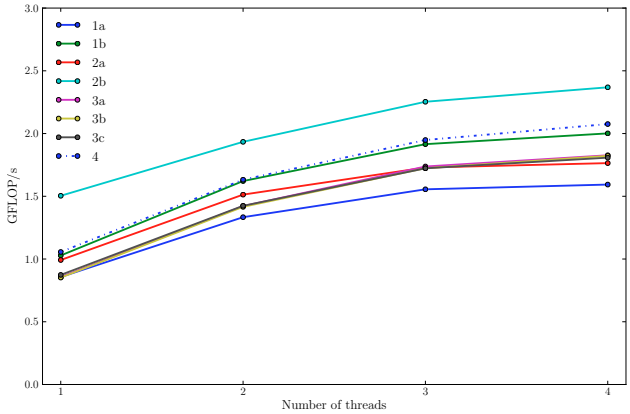

FIG. 5.5. *MKL SpMV GFlops per matrix for a varying number of threads.*

gesting the that underlying implementation is sequential. Indeed, our own sequential implementation of sparse matrix transpose offers equivalent performance.

The Galerkin Product results represent the cost of constructing the coarse-grid matrix by performing two sparse matrix-matrix multiplication operations. Whereas the cost of other components, such as BLAS algorithms, is directly proportional to the input size, the cost of sparse matrix-matrix multiplication is dependent on the specific sparsity patterns of the two matrices and can differ between inputs of the same size. Our sparse matrix-matrix multiplication algorithm achieves better performance than MKL in all cases considered, with an average speedup of 1.65. As with the transpose, MKL's sparse matrix-matrix multiplication functionality extracts no observable benefit from multithreading and offers inferior performance to our own sequential implementation based on the SMMP method [3]. As a result, the CPU performance results in the remainder of this section only use MKL for sparse matrix-vector multiplication.

5.4. Setup Phase Performance. Section 5.3 demonstrates that GPU implementations of the essential sparse matrix operations achieve tangible speedups over a competent CPU implementations. In this section we analyze the performance of the whole multigrid setup phase (cf. Algorith 1).

Figure 5.6 shows a breakdown of the parallel setup phase into the individual components. The figure identifies several intensive calculations in the algorithm, namely aggregation and the Galerkin product. As anticipated, the ability to express the strength-of-connection, tentative prolongator construction, and prolongator smoothing functions in terms of parallel primitives leads to a efficient execution on the GPU. Moreover, the figure also shows that matrix conversion is a relatively low cost, enabling the use of SpMV-optimized matrix formats in the subsequent cycling phase.

The Galerkin product, which is implemented with two sparse matrix-matrix multiplications $\mathtt{R} * (\mathtt{A} * \mathtt{P})$, is the most expensive step in all cases considered. The relative cost of the Galerkin product is slightly higher in the 3D matrix examples (2a, 2b, and 4) compared to the 2D cases.

In Table 5.7, as a baseline, we also include timings from the ML package of Trilinos on the same hardware; this validates that our CPU approach is in-fact a competitive implementation. The GPU implementation is faster in all eight examples, with an average speedup of 1.89 over the CPU reference.

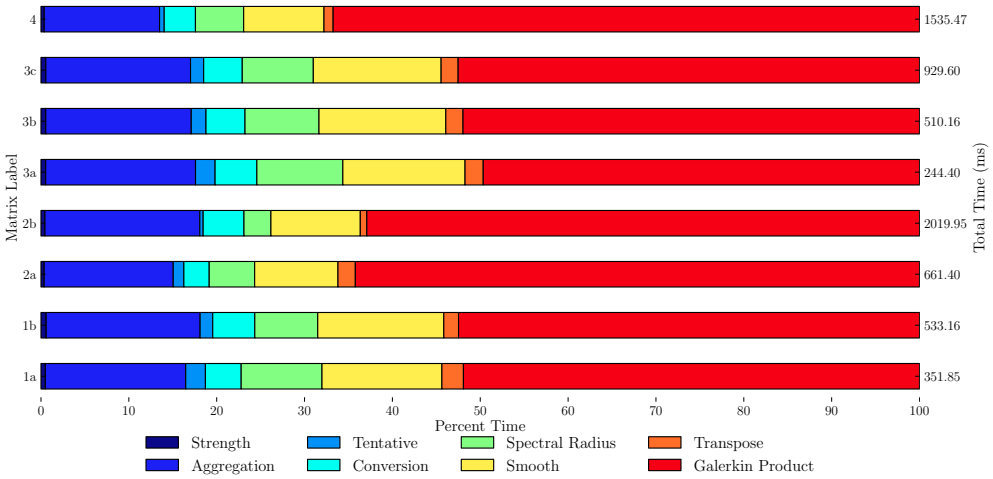

FIG. 5.6. *Relative time on the GPU of each setup phase component on the finest grid level.*

Matrix	CPU	GPU	Speedup	Trilinos-ML
1a. 2D FD	892	518	1.72	2040
1b. 2D FE	1133	649	1.75	2298
2a. 3D FD	1639	944	1.74	2906
2b. 3D FE	2845	2124	1.34	4420
3a. 2D FE	657	335	1.96	1324
3b. 2D FE	1484	648	2.29	2785
3c. 2D FE	2901	1151	2.52	5236
4. 3D FE	3157	1726	1.83	4967

TABLE 5.7
AMG setup time for all components (ms).

5.5. Solve Phase Performance. In this section we present the results of our AMG preconditioned CG solver for each matrix in our test suite. The AMG preconditioner uses a single iteration of weighted Jacobi at the of pre- and post-smoothing steps of the V-cycle. The Trilinos data was collected using the default uncoupled aggregation method with a tolerance of 10^{-12} and a maximum of 500 iterations.

Figure 5.8 shows that the main cost of the cycling is in pre- and post-smoothing, as desired. Since smoothing is of the form $x \leftarrow x + D^{-1}r$, pre-smoothing is minimal and the cost of constructing the residual is represented in $\mathtt{A} * \mathtt{x}$. This also shows that each major component of the cycle is comparable to a matrix-vector multiply. Furthermore, Figure 5.9 also shows that the entire AMG cycling is on the order of several matrix vector products, $\mathtt{A} * \mathtt{p}$.

Table 5.10 presents the results of the solve phase performance of our CPU and GPU implementations in addition to results from Trilinos ML. The increase in the number of iterations required by the GPU reflects the fact that, on structured problems, sequential aggregation method fortuitously selects square-shaped aggregates, while the parallel aggregation method based on MIS(2) creates irregularly shaped aggregates. Despite the increase in the number of iterations performed on the device, the

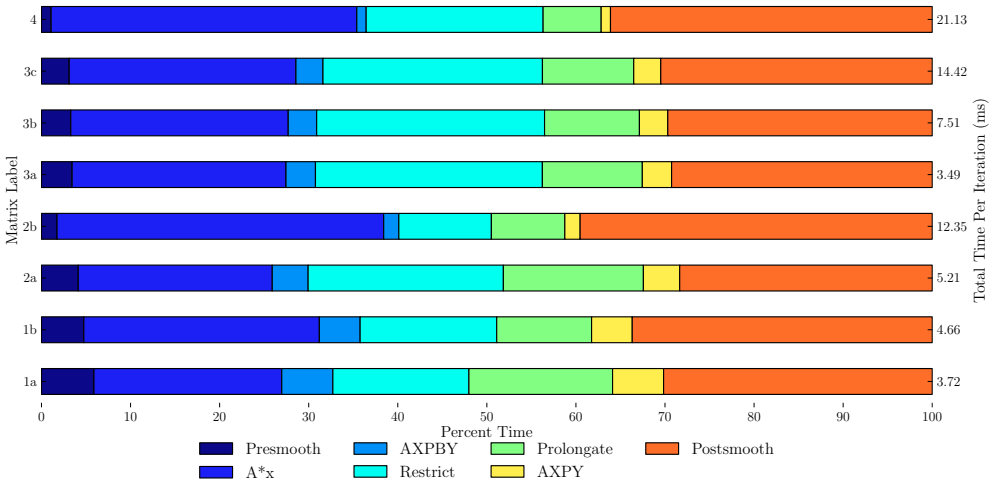

FIG. 5.8. *GPU V-cycle time breakdown on the finest grid level.*

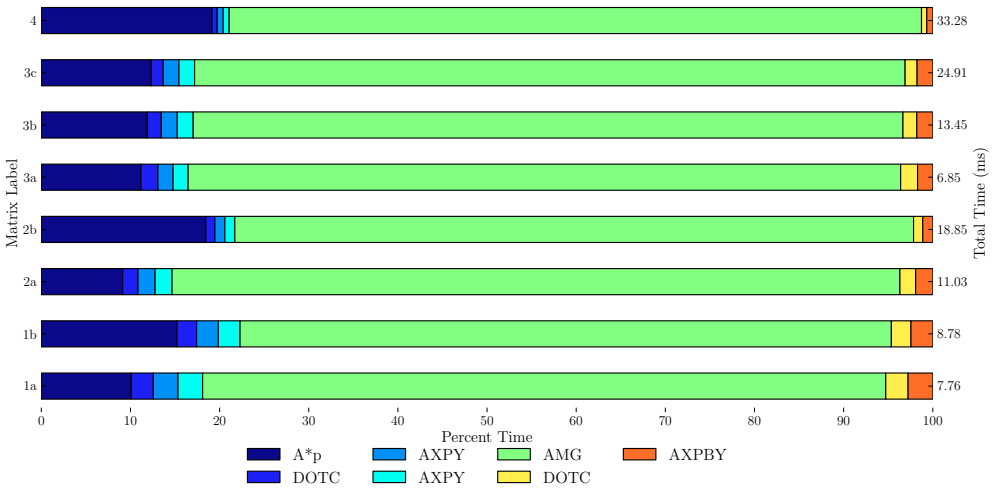

FIG. 5.9. *Total preconditioned solver time breakdown time on the GPU.*

increased SpMV performance using the HYB format allows the GPU to outperform the CPU in all cases.

Results on solve phase performance per level are presented in Figure 5.11. In each case, the same multigrid hierarchy is used on both the CPU and GPU to isolate the performance of the cycling phase. The results demonstrate that the GPU is noticeably faster than the CPU on the first two or three grid levels, which represents the vast majority of the computation in a multigrid V-cycle. The coarsest grid levels are processed more rapidly by the CPU, owing to communication latency and the fact that the amount of parallelism on the smaller grids is insufficient to saturate the GPU. However, since the absolute time spent on such grids is small, the GPU cycle is faster overall. However, the same is not necessarily true of other multigrid cycles, such as the F- or W-cycle, which visit coarse grids more frequently than fine grids.

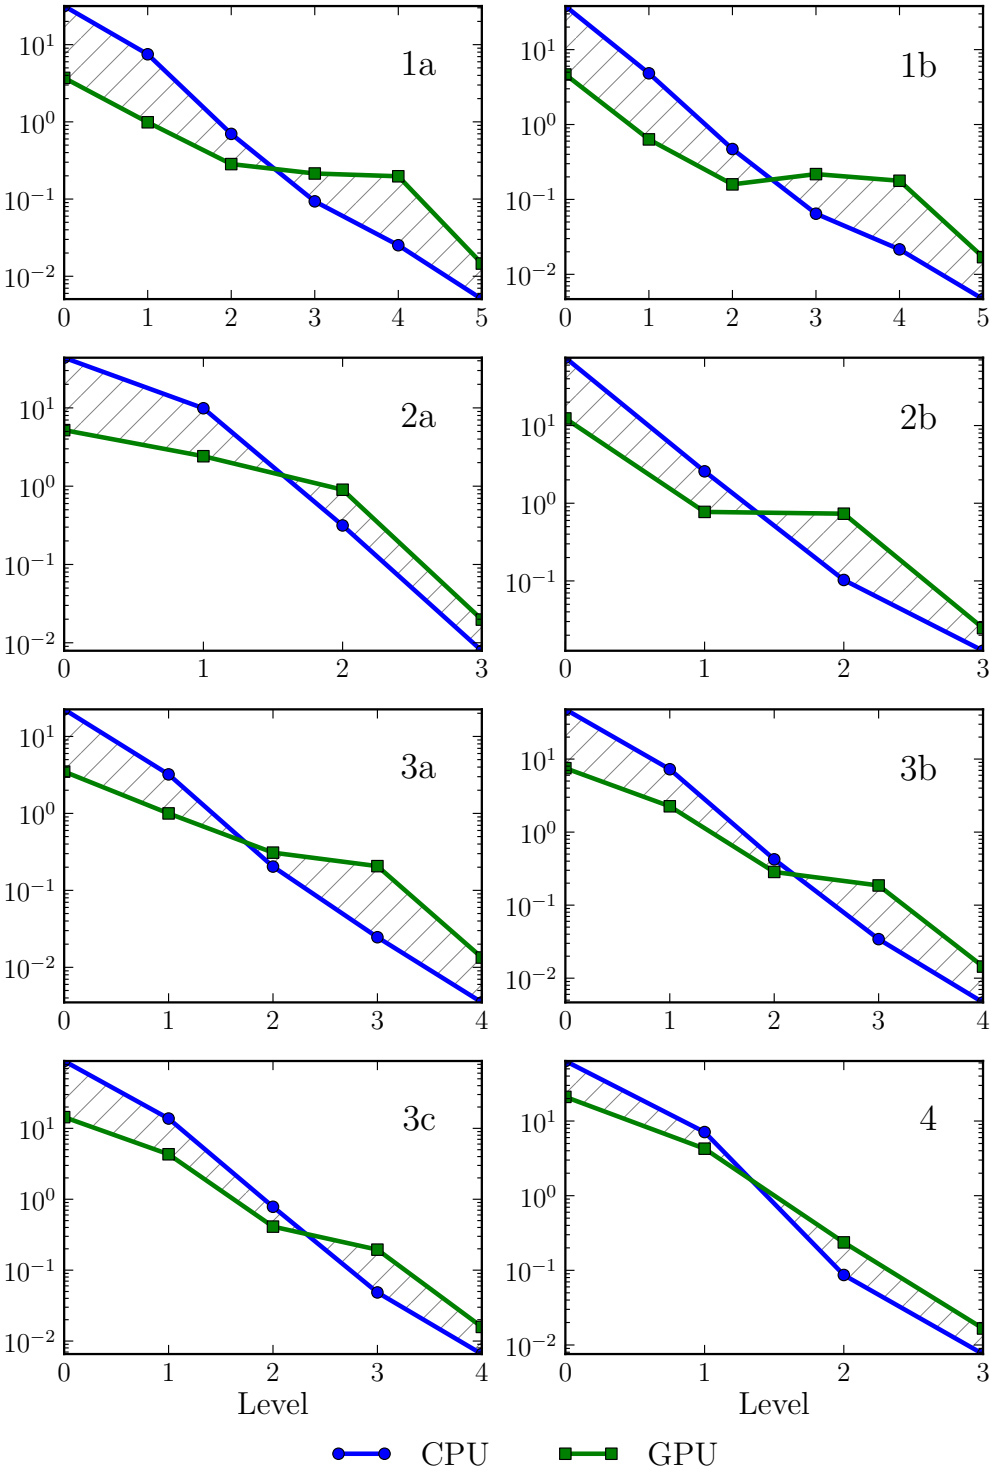

FIG. 5.11. *CPU vs GPU solve phase performance per level.*

Matrix	CPU		GPU		**Speedup**	ML	
	time	it.	time	it.	(per it.)	time	it.
1a. 2D FD	1221	20	423	51	**7.66**	14190	33
1b. 2D FE	1097	16	461	46	**7.52**	10590	22
2a. 3D FD	1760	23	295	27	**6.76**	14800	31
2b. 3D FE	1683	14	482	24	**5.98**	13840	20
3a. 2D FE	1534	42	286	49	**5.40**	14020	53
3b. 2D FE	3704	47	633	54	**5.77**	34410	68
3c. 2D FE	7804	53	1424	65	**5.75**	44530	65
4. 3D FE	4369	43	1498	50	**2.96**	28380	47

TABLE 5.10

AMG-PCG GPU solve time (ms), number of solver iterations and per-iteration speedup.

5.6. Scalability. One important feature of a successful multigrid method is the ability to scale linearly with matrix size n. Consequently, we expect our algorithm to also exhibit this scaling in observed setup and solve times. Here we consider a Example 1*a* above — i.e., the 2D structured Poisson problem — where algebraic multigrid is known to scale linearly. The problem size is scaled equally in each coordinate direction, and we measure the wall-clock time of both the setup and solve phase.

The results are highlighted in Figures 5.12 and 5.13. As shown in Figures 5.12a and 5.13a, we observe sub-linear scaling before the GPU is saturated, and the expected linear scaling after the problem size is sufficiently large. The fortuitous sublinear growth is consistent with other algorithms that utilize primitives on the GPU — e.g., see [30]. Indeed, as depicted in Figures 5.12b and 5.13b, if we measure the growth rate r of time dependence $\mathcal{O}(n^r)$ over a window of five samples, the scaling settles at a linear relationship.

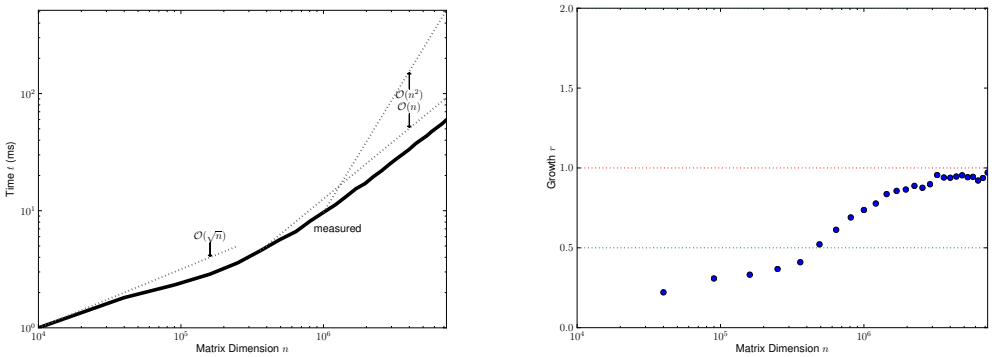

(a) Measured setup time in comparison to $\mathcal{O}(\sqrt{n})$, $\mathcal{O}(n)$, $\mathcal{O}(n^2)$.

(b) Approximate rate of growth r in $\mathcal{O}(n^r)$.

FIG. 5.12. *2D structured setup scalability test.*

6. Conclusions. We have demonstrated the first implementation of AMG that exposes fine-grained parallelism at all stages of the setup and cycling phases. In particular, we have described highly-parallel methods for sparse matrix-matrix mul-

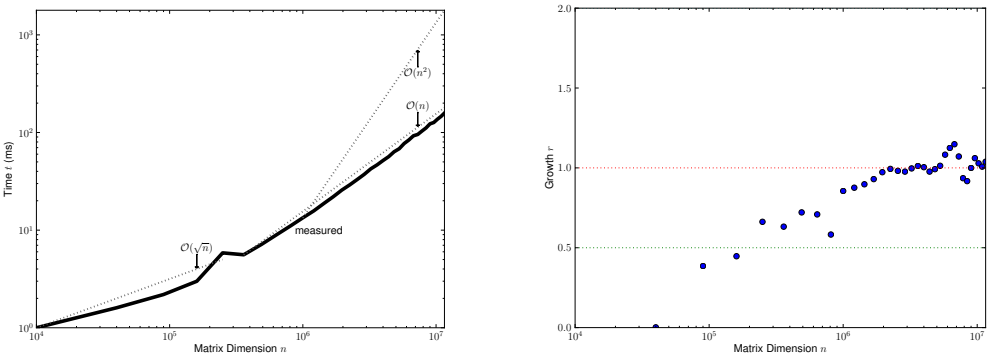

(a) Measured solve time in comparison to $\mathcal{O}(\sqrt{n})$, $\mathcal{O}(n)$, $\mathcal{O}(n^2)$.

(b) Approximate rate of growth r in $\mathcal{O}(n^r)$.

FIG. 5.13. *2D structured solve scalability test.*

tiplication and aggregation. Furthermore, we have shown that parallel primitives are a viable substrate to describing complex sparse matrix operations and targeting the GPU.

Building upon earlier work in sparse matrix-vector multiplication, we have demonstrated meaningful speedup in both the setup and cycling phases of the AMG solver on structured and unstructured problems. Whereas CPU implementations of major components such as sparse matrix-vector multiplication, sparse matrix-matrix multiplication, and sparse matrix transposition, derive little or no benefit from multithreading, our implementations leverage scalable parallel primitives.

Appendix. Distance-k Maximal Independent Sets.

In this section we describe an efficient parallel algorithm for computing distance-k maximal independent sets, denoted MIS(k) and defined in Definition 3.1. We begin with a discussion of the standard distance-1 maximal independent set — i.e., MIS(1) — and then detail the generalization to arbitrary distances. Our primary interest is in computing a MIS(2) to be used in the parallel aggregation scheme discussed in Section 3.2.

Computing a distance-1 maximal independent set is straightforward in a serial setting, as shown by Algorithm 4. The algorithm is greedy and iterates over nodes, labeling some as MIS nodes and their neighbors as non-MIS nodes. Specifically, all nodes are initially candidates for membership in the maximal independent set s and labeled (with value 0) as undecided. When a candidate node is encountered it is labeled (with value 1) as a member of the MIS, and all candidate neighbors of the MIS node are labeled (with value -1) as being removed from the MIS. Upon completion, the candidate set is empty and all nodes are labeled with either a 0 or 1.

Computing maximal independent sets in parallel is challenging, the but several methods exist. With $k = 2$, our parallel version in Algorithm 5 can be considered a variant of Luby's method [28] which has been employed in many codes such as ParMETIS [25]. A common characteristic of such schemes is the use of randomization to select independent set nodes in parallel.

As with the serial method, all nodes are initially labeled (with a 0) as a candidate for membership in the MIS s. Additionally, each node is assigned a random value in the array v. The purpose of the random value is to create disorder in the selection of

Algorithm 4: `MIS_serial`

parameters: A, $n \times n$ sparse matrix
return: s, independent set

$I = \{0, \ldots, n\}$ {initial candidate index set}
$s \leftarrow 0$ {initialize to undecided}

for $i \in I$ {for each candidate}
  if $s_i = 0$ {if unmarked...}
    $s_i = 1$ {add to candidate set}
    for j *such that* $A_{ij} \neq 0$
      $s_j = -1$ {remove neighbors from candidate set}

$s = \{i : s_i = 1\}$ {return a list of MIS nodes}

nodes, allowing many nodes to be added to the independent set at once. Specifically, the random values represent the precedence in which nodes are considered for membership in the independent set. In the serial method this precedence is implicit in the graph ordering. Figure A.1 illustrates a two-dimensional graph with random values values drawn from integers in $[0, n = 36)$.

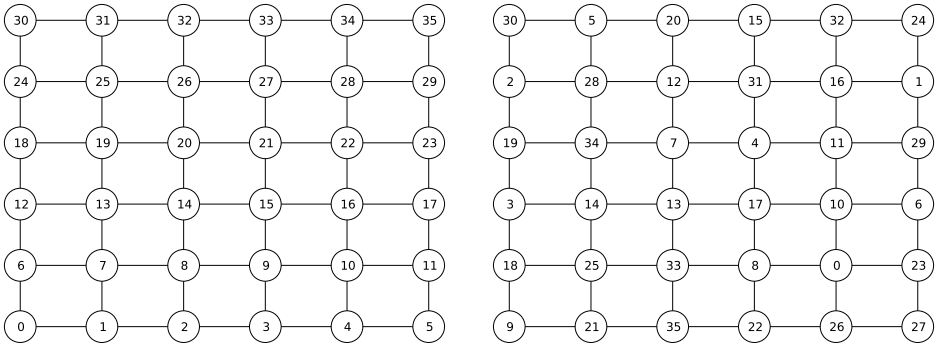

(a) Natural enumeration. (b) Random enumeration.

FIG. A.1. *A structured graph with a natural order of nodes and a randomized enumeration.*

The algorithm iterates until all nodes have been labeled with a -1 or 1, classifying them as a non-MIS node or MIS node respectively. In each iteration an array T of 3-tuples is created, tying together the node state — i.e., -1, 0, or 1 — the node's random value, and the linear index of the node. In a second phase, the nodes compute, in parallel, the maximum tuple among the neighbors. Given two tuples $t_i = (s_i, v_i, i)$ and $t_j = (s_j, v_j, j)$ the maximum is determined by a lexicographical ordering of the tuples. For example the tuple $(2, 5, 10)$ is larger than $(1, 10, 20)$ but smaller than $(2, 6, 0)$. Note that this ordering ensures that MIS nodes have a priority over candidate nodes and that candidate nodes have priority over non-MIS nodes. In the third phase, the candidate node states are updated based on the results of the second phase. Candidate nodes that are the local maximum — i.e., $s_{max} = i$ — are added to the independent set while those with an existing MIS neighbor — i.e., $s_{max} = 2$ — are removed from candidacy. Since it is impossible for two neighboring nodes to be local maxima, the

Algorithm 5: `MIS_parallel`

```
parameters: A   n×n sparse matrix
            k   edge distance
return:     s,  independent set

I = {0, ..., numrows(A)}

s ← 0                                          {initialize state as undecided}
v ←random                                      {initialize value}

while {i ∈ I : s_i = 1} ≠ ∅
    for i ∈ I
        T_i ← (s_i, v_i, i)                    {set tuple (state,value,index)}

1   for r = 1, ..., k                          {propagate distance k}
        for i ∈ I
            t ← T_i
            for j such that A_ij ≠ 0
                t ← max(t, T_j)                {maximal tuple among neighbors}
            T̂_i ← t
        T = T̂

2   for i ∈ I
        (s_max, v_max, i_max) ← T_i
        if s_i = 0                             {if unmarked...}
            if i_max = i                       {if maximal...}
                s_i ← 1                        {add to set}
            else if s_max = 1                  {otherwise...}
                s_i ← −1                       {remove from set}

s = {i : s_i = 1}                              {return a list of MIS nodes}
```

selected nodes are independent by construction. Furthermore, the correctness of the algorithm does not depend on the random values. Indeed, if all the random values are 0, the algorithm degenerates into the serial algorithm since the third component of the tuple, the node index, establishes precedence among neighbors with the same random value. In each iteration of the algorithm at least one candidate node's state is changed, so termination is assured. Figure A.2 illustrates the classification of nodes during six iterations of the parallel algorithm.

Generalizing the parallel distance-1 maximal independent set algorithm to compute MIS(k) for arbitrary k can be accomplished in several ways. One solution is to compute the k-th power of A as an explicit matrix $A^K = A * A * \ldots A$ using sparse matrix-matrix multiplication, and then to apply the standard distance-1 maximal independent set algorithm. However, computing A^k explicitly is computationally expensive and the storage for A^k grows rapidly with k. An alternative approach,

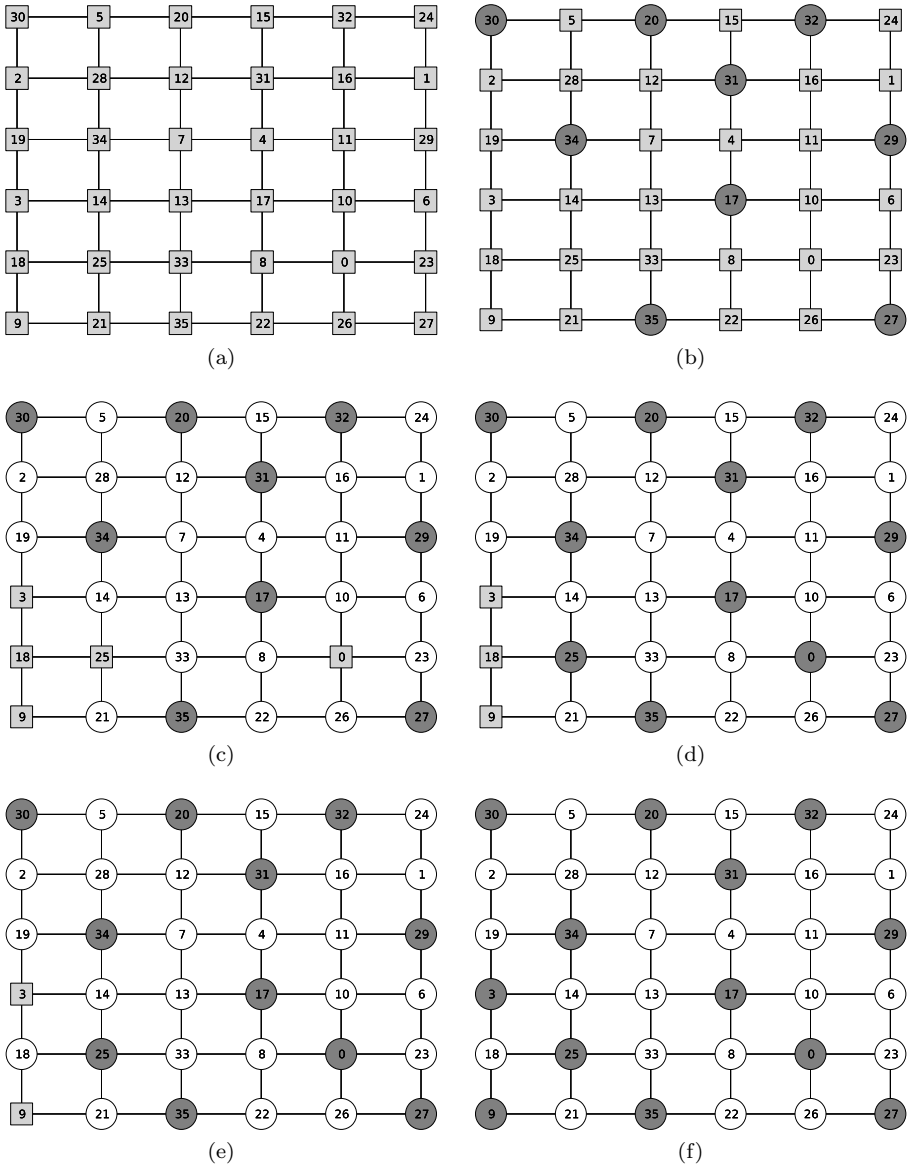

FIG. A.2. *Parallel MIS construction for* $k = 1$.

which is generally superior, uses $k > 1$ in Algorithm 5 so that maximum tuples are propagated k times using A, which has the same effect as one step of A^k. Line 1 of Algorithm 5 illustrates this scheme, with an additional optimization in the form of a second state-update pass. The second update pass on Line 2 exploits the fact that many nodes can be immediately classified based on the results of the first pass without another iteration. Specifically, nodes whose maximum neighbor was classified as an independent set node — i.e., $s_{n_{max}} = 0$ — can be safely removed from candidacy. This optimization generally reduces the number of outer iterations by anticipating and efficiently applying the effect of the *next* iteration. An example of distance-2 MIS is depicted in Figure A.3.

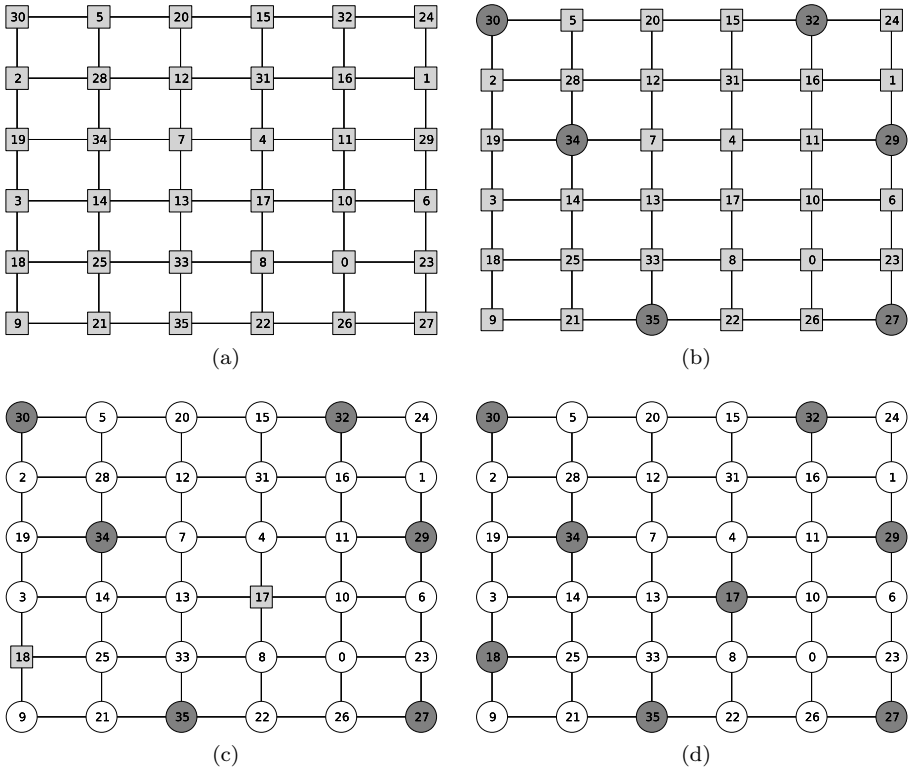

FIG. A.3. *Parallel MIS construction for* $k = 2$.

We note that this optimization does not capture all cases where a node is removed from candidacy in the next iteration. Nevertheless, the second pass generally a worthwhile optimization.

In our implementation of Algorithm 5 the random values are produced by a hash function, $v_i = \texttt{hash}(i)$, where `hash` is a simple integer hash function. Although not a source of high quality random numbers, the resulting values are adequate for our purpose. More sophisticated hash-based random number generators are discussed in [36, 40].

Figure A.4 shows the results of an empirical test of our MIS(2)-based aggregation algorithm. We compare the convergence rate of a solver constructed with the standard serial aggregation algorithm against the convergence rate of a solver constructed with our parallel aggregation algorithm when applied to a small 2D Poisson problem. The histogram represents the distribution of convergence rates across 10,000 trials of each method (both serial and parallel). In each trial the rows and columns of the matrix are permuted randomly using a high-quality pseudo-random number generator and a two-level hierarchy is constructed from using the serial and parallel aggregation schemes on the permuted matrix. Randomly permuting the matrix has the effect of randomizing the order in which the sequential aggregation algorithm visits nodes.

The close agreement of the two distributions offers empirical evidence that our hash-based randomization method is adequate for the purpose of computing aggregates in parallel. Indeed, the average (mean) convergence rates of the serial and parallel methods are 0.7582 and 0.7580 respectively.

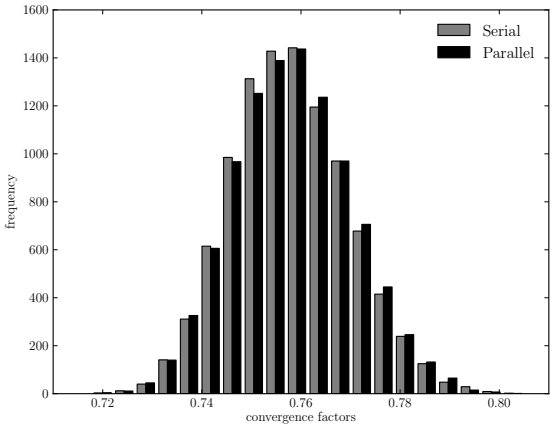

FIG. A.4. *Distribution of convergence factors for serial and parallel aggregation methods.*

REFERENCES

[1] MARK ADAMS, MARIAN BREZINA, JONATHAN HU, AND RAY TUMINARO, *Parallel multigrid smoothing: polynomial versus Gauss-Seidel*, J. Comput. Phys., 188 (2003), pp. 593–610.

[2] A.H. BAKER, T. GAMBLIN, M. SCHULZ, AND U.M. YANG, *Challenges of scaling algebraic multigrid across modern multicore architectures*, in Proceedings of the 25th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2011), 2011.

[3] R.E. BANK AND C.C. DOUGLAS, *Sparse matrix multiplication package (SMMP)*, Advances in Computational Mathematics, 1 (1993), pp. 127–137.

[4] MUTHU MANIKANDAN BASKARAN AND RAJESH BORDAWEKAR, *Optimizing sparse matrix-vector multiplication on GPUs*, IBM Research Report RC24704, IBM, Apr. 2009.

[5] NATHAN BELL AND MICHAEL GARLAND, *Efficient sparse matrix-vector multiplication on CUDA*, NVIDIA Technical Report NVR-2008-004, NVIDIA Corporation, Dec. 2008.

[6] ———, *CUSP : Generic parallel algorithms for sparse matrix and graph computations*. `http://code.google.com/p/cusp-library/`, 2009-.

[7] NATHAN BELL AND MICHAEL GARLAND, *Implementing sparse matrix-vector multiplication on throughput-oriented processors*, in Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, SC '09, New York, NY, USA, 2009, ACM, pp. 18:1–18:11.

[8] GUY E. BLELLOCH, *Vector models for data-parallel computing*, MIT Press Cambridge, Mass, 1990.

[9] JEFF BOLZ, IAN FARMER, EITAN GRINSPUN, AND PETER SCHRÖODER, *Sparse matrix solvers on the GPU: conjugate gradients and multigrid*, ACM Trans. Graph., 22 (2003), pp. 917–924.

[10] E. CHOW, R.D. FALGOUT, J.J. HU, R.S. TUMINARO, AND U.M. YANG, *A survey of parallelization techniques for multigrid solvers*, Parallel processing for scientific computing, (2006), pp. 179–202.

[11] M. CHRISTEN, O. SCHENK, AND H. BURKHART, *General-purpose sparse matrix building blocks using the NVIDIA CUDA technology platform*, in First Workshop on General Purpose Processing on Graphics Processing Units, Northeastern Univ., Boston, 2007.

[12] A.J. CLEARY, R.D. FALGOUT, V.E. HENSON, J.E. JONES, T.A. MANTEUFFEL, S.F. MCCORMICK, G.N. MIRANDA, AND J.W. RUGE, *Robustness and scalability of algebraic multigrid*, SIAM Journal on Scientific Computing, 21 (2000), pp. 1886–1908.

[13] JONATHAN M. COHEN AND JEROEN MOLEMAKE, *A fast double precision CFD code using CUDA*, in 21st International Conference on Parallel Computational Fluid Dynamics (ParCFD2009), 2009.

[14] NVIDIA CORPORATION, *Cublas library*, May 2010. version 3.1.

[15] JEFFREY DEAN AND SANJAY GHEMAWAT, *MapReduce: Simplified data processing on large clusters*, Commun. ACM, 51 (2008), pp. 107–113.

[16] E.F. DAZEVEDO, M.R. FAHEY, AND R.T. MILLS, *Vectorized sparse matrix multiply for com-*

pressed row storage format, in International Conference on Computational Science (ICCS), Springer, 2005, pp. 99–106.

[17] Michael Garland and David B. Kirk, *Understanding throughput-oriented architectures*, Commun. ACM, 53 (2010), pp. 58–66.

[18] M.W. Gee, C.M. Siefert, J.J. Hu, R.S. Tuminaro, and M.G. Sala, *ML 5.0 smoothed aggregation user's guide*, Tech. Report SAND2006-2649, Sandia National Laboratories, 2006.

[19] Dominik Göddeke, Robert Strzodka, Jamaludin Mohd-Yusof, Patrick S. McCormick, Hilmar Wobker, Christian Becker, and Stefan Turek, *Using GPUs to improve multigrid solver performance on a cluster*, International Journal of Computational Science and Engineering, 4 (2008), pp. 36–55.

[20] Nolan Goodnight, Cliff Woolley, Gregory Lewin, David Luebke, and Greg Humphreys, *A multigrid solver for boundary value problems using programmable graphics hardware*, in HWWS '03: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware, Aire-la-Ville, Switzerland, Switzerland, 2003, Eurographics Association, pp. 102–111.

[21] Fred G. Gustavson, *Two fast algorithms for sparse matrices: Multiplication and permuted transposition*, ACM Trans. Math. Softw., 4 (1978), pp. 250–269.

[22] G. Haase, M. Liebmann, G. Plank, and C. Douglas, *Parallel algebraic multigrid on general purpose gpus*, in 3rd Austrian Grid Symposium, J. Volkert et al, ed., 2010, pp. 28–37.

[23] Van Emden Henson and Ulrike Meier Yang, *Boomeramg: A parallel algebraic multigrid solver and preconditioner*, Applied Numerical Mathematics, 41 (2002), pp. 155 – 177.

[24] Jared Hoberock and Nathan Bell, *Thrust: A parallel template library*, 2011. Version 1.4.0.

[25] George Karypis and Vipin Kumar, *Parallel multilevel k-way partitioning scheme for irregular graphs*, SIAM Review, 41 (1999), pp. 278–300.

[26] Michael Kazhdan and Hugues Hoppe, *Streaming multigrid for gradient-domain operations on large images*, in ACM SIGGRAPH 2008 papers, SIGGRAPH '08, New York, NY, USA, 2008, ACM, pp. 21:1–21:10.

[27] CL Lawson, RJ Hanson, DR Kincaid, and FT Krogh, *Basic Linear Algebra Subprograms for Fortran Usage*, ACM Transactions on Mathematical Software (TOMS), 5 (1979), pp. 308–323.

[28] Michael Luby, *A simple parallel algorithm for the maximal independent set problem*, SIAM J. Comput., 15 (1986), pp. 1036–1055.

[29] Duane Merrill and Andrew Grimshaw, *Parallel scan for stream architectures*, Tech. Report CS2009-14, University of Virginia, Department of Computer Science, Charlottesville, VA, USA, 2009.

[30] ———, *Revisiting sorting for gpgpu stream architectures*, Tech. Report CS2010-03, University of Virginia, Department of Computer Science, Charlottesville, VA, USA, 2010.

[31] NVIDIA Corporation, *NVIDIA CUDA Programming Guide*, May 2011. Version 4.0.

[32] J. W. Ruge and K. Stüben, *Algebraic multigrid*, in Multigrid methods, vol. 3 of Frontiers Appl. Math., SIAM, Philadelphia, PA, 1987, pp. 73–130.

[33] Shubhabrata Sengupta, Mark Harris, Yao Zhang, and John D. Owens, *Scan primitives for GPU computing*, in Graphics Hardware 2007, ACM, Aug. 2007, pp. 97–106.

[34] Markus Stürmer, Harald Köstler, and Ulrich Rüde, *How to optimize geometric multigrid methods on gpus*, in 15th Copper Mountain Conference on Multigrid Methods, March 2011.

[35] Ray S. Tuminaro and Charles Tong, *Parallel smoothed aggregation multigrid : Aggregation strategies on massively parallel machines*, SC Conference, 0 (2000), p. 5.

[36] S. Tzeng and L.Y. Wei, *Parallel white noise generation on a gpu via cryptographic hash*, in Proceedings of the 2008 symposium on Interactive 3D graphics and games, ACM, 2008, pp. 79–87.

[37] P. Vaněk, J. Mandel, and M. Brezina, *Algebraic Multigrid by Smoothed Aggregation for Second and Fourth Order Elliptic Problems*, Computing, 56 (1996), pp. 179–196.

[38] Richard Vuduc and Hyun-Jin Moon, *Fast Sparse Matrix-Vector Multiplication by Exploiting Variable Block Structure*, in High Performance Computing And Communications: First International Conference, HPCC 2005, Sorrento, Italy, September 21-23, 2005: Proceedings, Springer, 2005.

[39] S. Williams, L. Oliker, R. Vuduc, J. Shalf, K. Yelick, and J. Demmel, *Optimization of sparse matrix-vector multiplication on emerging multicore platforms*, in Proc. 2007 ACM/IEEE Conference on Supercomputing, 2007.

[40] F. Zafar, M. Olano, and A. Curtis, *Gpu random numbers via the tiny encryption algorithm*, in Proceedings of the Conference on High Performance Graphics, Eurographics Association, 2010, pp. 133–141.